Pricing for Efficient Traffic Exchange at IXPs

Md Ibrahim Ibne Alam, *Member, IEEE,* Elliot Anshelevich, *Member, ACM,*
Koushik Kar, *Senior Member, IEEE,* and Murat Yuksel, *Senior Member, IEEE*

Abstract—We analyze traffic exchange between Internet Service Providers (ISPs) at an Internet Exchange Point (IXP) as a non-cooperative game with ISPs as self-interested agents. Each ISP has the choice of exchanging traffic either using the shared IXP facilities, or outside the IXP – through their transit providers or private peering. We analyze the efficiency (social cost optimality) of the traffic exchange equilibrium at the IXP taking into consideration the congestion cost experienced by the ISPs at the IXP. To model both non-profit and for profit IXPs, we consider several cases, i) where the IXP does not charge any price to ISPs for the traffic exchanged (zero pricing), ii) when it charges a price that is proportional to the aggregate level of congestion at the IXP (proportional pricing), and iii) when it charges a constant price per unit traffic (constant pricing). Further, we also analyze the profit earned by the IXP under these pricing policies, under two different models of the congestion cost (delay) functions. Simulations conducted using data for actual IXPs obtained from PeeringDB demonstrate that the theoretical bounds derived for social cost and profit optimality at equilibrium (measured as the Price of Anarchy) are fairly tight, and correctly capture the performance trends against the variation of key model parameters. Further, the results show that for proportional pricing, there is an operating price range that attains near-optimal social cost and near-optimal IXP profit *simultaneously*. We also demonstrate - through both theoretical analysis and simulations - that as compared to zero and constant pricing policies, proportional pricing attains better tradeoff between social cost and IXP profit, and also results in a performance that is more robust to price variations.

Index Terms—Traffic exchange games, traffic pricing, traffic equilibrium.

1 INTRODUCTION

1.1 Background and Motivation

INTERNET service providers typically connect (mostly peer) with each other at Internet eXchange Points (IXPs). In most basic terms, an IXP is a data center with network switches through which ISPs form connections (peering relationships) to exchange traffic [2], [3]. In return, IXPs recover their operating costs by charging fees to each member/client ISP; and the operating costs of an IXP is largely determined by the cost of the infrastructure needed for traffic exchange [4]. A number of IXPs, especially in Europe, operate on a non-profit basis [5], whereas other IXPs, both in Europe and particularly US, operate for profit, e.g., Equinix [6]. In both cases, while IXPs provide the platform for ISPs to connect (peer) with each other, they play a passive role focused on infrastructure cost recovery or profit-making, and the peering decisions are determined bilaterally by the ISPs themselves. Nevertheless, ISPs at an IXP make these peering decisions taking into account the potential quality of service (bandwidth, delay etc.) improvements due to peering, the prices charged by the IXP, and comparing those with alternatives such as sending the traffic through their transit providers.

In recent years, transit prices per unit bandwidth have been steadily declining [7]. Despite falling transit costs, peering between ISPs has been on the rise, and content and access ISPs are increasingly getting into peering relationships [8], [9], a phenomenon known as the flattening of the Internet [3], [10], [11], [12]. It has been shown [13] that almost 80% of the IP addresses can be reached via public peering, and 20% of all the traffic traces go through IXPs. Peering between ISPs, which is typically settlement-free, can help bring contents closer to customers, resulting in lower delays and losses, and thus better Quality-of-Experience (QoE) for the end users (content consumers). Some of the recent literature has therefore argued that *paid peering* is necessary for overall stability and efficiency of inter-domain traffic [14], though its interplay with traditional settlement-free peering needs careful treatment [15].

The importance of IXP-driven innovation in shaping and redefining the ISP traffic exchange marketplace (which is already happening in Europe) has been highlighted in several recent studies such as [16], [17].

With recent developments in software-defined networking, more controls are enabled at IXPs in terms of managing flows [18], [19], maintaining peering relationships [20], [21], and practicing more scalable architectures [22], [23] in the face of growing amount of traffic passing through their datacenters. These innovations help in reducing peering costs and invite IXPs to a more active role in Internet peering.

The peering decisions are ultimately left up to the ISPs, which make these decisions based on not only the cost of using the IXP but also other costs (such as transit costs) that are not generally known to the IXP. Even though the IXPs play a passive role in the ISP connectivity by providing peering facilities for a price, the pricing policy applied by the IXP to facilitate this traffic exchange needs to be designed carefully for making the peering relationships and traffic flows between the ISPs stable and efficient. The price paid for traffic exchange at an IXP can include costs that are not shared with other ISPs, such as cost of transport to the IXP and co-location cost [24], [25]. In this paper, we only focus on pricing of the shared resource – the public switch at the IXP

This is an extended version of the work in [1].
Md Ibrahim Ibne Alam, Elliot Anshelevich and Koushik Kar are with Rensselaer Polytechnic Institute; Emails: {alamm2, eanshel@cs, kark}@rpi.edu.
Murat Yuksel is with the University of Central Florida; Email: Murat Yuksel Murat.Yuksel@ucf.edu.
Manuscript received xxx xx, xxxx; revised xxx xx, xxxx.

where ISPs exchange traffic with each other. The price levied on an ISP for the use of this shared resource has to account for the overall congestion (or in other words the traffic sent by other ISPs) necessitating a game theoretic treatment. In current practice, the price paid by an ISP for sharing an IXP switch depends only on the port capacity it purchases and not on the port capacity of the other ISPs [24], [26]. Therefore, the current pricing policy, at least in the short term, do not directly factor in the overall congestion level of the switch measured in terms of the total port capacity purchased or the total traffic exchanged through the switch. However, congestion-insensitive pricing can lead to poor traffic exchange solutions at equilibrium. This motivates us to analyze congestion-dependent IXP pricing policies, which, when designed carefully, can lead to efficient traffic exchange as we show in this paper.

Since the ISPs can be expected to make peering decisions in self-interest, the traffic pricing policy can be modeled as a non-cooperative game, where the quality of the equilibrium (resulting from an IXP's pricing policy) must be measured in terms of the efficiency of the equilibrium traffic flow. The consideration of profit-making IXPs introduces another dimension to this game, and the IXP can now be considered as a selfish player that would ideally like to set prices to maximize its profit. Not suprisingly, it can be shown that pricing policies that minimize social cost (i.e., maximize traffic flow efficiency between the ISPs connected to the IXP) can result in poor IXP profit, and vice versa. The question that naturally arises is whether there exists a pricing policy that can attain a good trade-off between social cost and IXP profit, i.e., it will result in a solution that is guaranteed to be within a small factor of the optimum in terms of both objectives (social cost as well as IXP profit).

1.2 Contribution of this Work

This paper investigates how the pricing policy at an IXP impacts the efficiency of the peering relationships that form between the ISPs as a result of that policy. We recognize that ISPs have a choice in terms of how they exchange traffic between each other: either through public peering at the IXP (using a shared switch); or outside of the IXP, through their transit providers or private peering. We define the traffic exchange problem between ISPs (at an IXP) as a non-cooperative game between ISPs (selfish agents), where each pair of ISPs have a certain pre-determined amount of traffic to exchange, and the strategic decision involves determining whether to send this traffic through the IXP or through the external routing option. We not only consider the traffic flow efficiency (captured by the social cost) at equilibrium, but also the profit earned by the IXP at equilibrium, and how the social cost and profit can be balanced carefully. We consider a *proportional pricing* policy, where the price charged by the IXP per unit traffic is proportional to the aggregate level of congestion at the IXP (shared switch). The benefit of proportional pricing – as discussed later in more detail – is that it can be implemented with only aggregate load information at the shared switch used for public peering. In other words, choosing a good pricing does not require the IXP to know the transit options or other details about the participating ISPs. Further, as we establish in this work, proportional pricing can also attain equilibrium solutions that are close to optimal both in terms of traffic flow efficiency (social or system objective) and IXP profit (market-maker's objective), for a range of choices of the proportionality constant. We also analyze the performance benefits of proportional pricing over zero pricing where the IXP does not charge any price at all.[1] Furthermore, we also analyze a constant pricing policy where the IXP chooses to charge a constant price per unit traffic, and comment on the implementation and performance advantages of proportional pricing over a constant pricing policy.

More specifically, this work makes the following contributions. First, we characterize the pricing policy that is economically efficient, i.e., attains the socially optimal traffic exchange solution. We then characterize the traffic flow efficiency (i.e., social cost optimality) of the equilibrium solutions – measured by the Price of Anarchy (PoA) of the system – under *proportional pricing*, where IXPs charge traffic a per-unit price that is proportional to the average level of congestion experienced by traffic in the public switch operated by the IXP. The quality of service experienced by the ISPs also suffers due to this congestion, and is taken into account as well. The PoA for *zero pricing*, where the traffic through the IXP only experiences a congestion cost, but no additional price is charged by the IXP, also follows from this result as a special case. In both of these pricing schemes, an IXP, in order to choose good prices, does *not* need to have detailed information about the costs of the ISPs' external routing options. This is in contrast to a *constant pricing* policy – which we also analyze theoretically – where each ISP pays a constant per-unit traffic price, correct calibration of which requires knowing the ISPs' external routing costs. We quantify the PoA under these pricing systems using two broad classes of delay functions (polynomial delay and queuing delay functions), and also discuss how the PoA results generalize when the external routing costs between ISPs are asymmetric. We then analyze the profit earned by the IXP (as compared to the optimal profit) under the pricing policies, for both polynomial and queuing delay functions. Finally, we simulate real-world scenarios of public peering in IXPs, and compare the actual performance values obtained against their respective theoretical bounds for a wide range of model parameters.

Our theoretical results show that the PoA (i.e., the measure of social optimality) under proportional pricing evaluates to a small constant for a wide range of model parameters. Numerical simulations using data from 28 large IXP locations in the US demonstrate that, under a wide range of operating conditions, the equilibrium solutions have an efficiency that is usually within a factor of 2 of the optimum. Further, the results show that for proportional pricing there is a price range that is desirable from both social cost and IXP profit perspectives, i.e., the social cost is close to the optimum, and at the same time the profit earned by the IXP is also quite close to the maximum possible. This is particularly significant since the proportional pricing policy that attains near-optimal social cost and IXP profit

1. Even with zero pricing, ISPs still face congestion costs at the shared switch and must determine its routing/peering choice taking into account how that cost compares with its alternate routing option.

only requires knowledge of the average congestion level at the public switch (which in turn depends on the volume of traffic exchanged through the IXP), and is therefore known to the IXP. In contrast, we argue that when a constant pricing policy is used, the choice of this constant per-unit traffic price that attains good PoA requires knowledge of the external routing costs of the ISPs which may not be practical. Further, unlike proportional pricing, both social cost and IXP profit are particularly sensitive to the choice of the constant price, and in general it is difficult to find a constant price that attains a good tradeoff between social cost and IXP profit across a wide range of model parameters.

1.3 Comparison with Prior Work

This work is part of an enormous line of work on network routing and formation by self-interested agents. Perhaps the most well-known line of work in this area is that of the foundations of selfish routing and congestion games, see Chapters 17-19 in [27], as well as [28] and many subsequent results. In this well-studied setting, agents choose routes on which to send their traffic, and much of the existing work studies Nash equilibrium (or its related notion of Wardrop equilibrium) in this setting. When modeling the behavior of ISPs and IXP pricing, however, such notions of equilibrium are no longer appropriate. This is due to the fact that when deciding to send traffic through an IXP, *both* participating ISPs must agree to this traffic through a peering contract: it is no longer a unilateral decision by a single agent. Nash equilibrium trivially exists: a solution where no one sends any traffic, for example, is a trvial (terrible) Nash equilibrium. Because of this, standard techniques such as congestion games or the potential function method [29] can no longer be directly applied to our setting, and a different notion of equilibrium is needed.

Because of the above observations, our game-theoretic model and analysis are inspired by a prior line of work on network formation games (introduced in [30]), where the stability of networks was modeled and analyzed when two nodes can only build links mutually but can sever links individually. These types of network formation games and their extensions have been studied extensively for different settings (e.g., [31], [32], [33], [34], [35], [36], [37], [38], [39]). Unlike these prior studies, in our model the cost of forming these (peering) connections is not fixed, but depends on both the congestion (measured by the total number of connections already formed), and the prices charged by the planner (IXP). In this sense, our work is a generalization of the previous work on network formation games, as it includes a central planner (the IXP), who can greatly affect the quality of outcomes by choosing different pricing schemes. There are many prior works on pricing network services and traffic ([40], [41], [42], [43], [44], [45], to list a few), but these models do not consider an IXP setting and hence are not directly related to ours.

Our work is most related to the model in [46], but differs from this existing work in several important aspects. First, while [46] considers the question of how the operational cost of the IXP should be shared among the ISPs (which is more representative of non-profit IXP operations), in our model the IXP directly charges the ISPs for their traffic. Secondly, we also analyze the profit earned by IXPs at equilibrium, and the trade-offs between social cost and IXP profit. Finally, our work also models congestion cost at the IXP, considers asymmetric external routing costs and paid peering, and evaluates social optimality and IXP profit for the pricing policies through extensive simulations.

1.4 Paper Organization

This paper is structured as follows. Section 2 describes the system model and derives some equilibrium properties that is useful in further analysis. Section 3 analyzes the efficiency of the equilibrium solution under proportional and constant pricing under two different congestion cost models. Section 4 analyzes IXP profit under the same pricing and congestion cost models. We present simulation results in Section 5, and conclude in Section 6.

2 System Model and Properties

2.1 Game-Theoretic Model

We consider an IXP, and a set N of ISPs (agents in our game-theoretic model) that are involved in traffic exchange through a public switch offered by the IXP. An ISP pair (i, j) has a total traffic demand of B_{ij} between themselves; part of this traffic, y_{ij}, is routed through the public switch, while the rest is sent *externally*. The traffic sent externally (i.e., outside of the public switch) is typically done in one of two ways: (a) through private peering between ISPs i and j;[2] (b) through the use of the ISPs' transit service providers. Generally, when an ISP joins an IXP, it gains access to other ISPs, but this does not necessarily mean it will exchange traffic (or have a peering relationship) with every other ISP present at that IXP. The traffic that is exchanged through the public switch incurs a congestion cost of $d(y)$ per unit traffic, which depends on the total traffic y sent through the switch. This congestion cost will typically be reflected in terms of average delay experienced by the traffic (and therefore we will sometimes use the terms 'congestion cost' and 'delay' interchangeably); however, $d(y)$ could also represent other Quality-of-Service (QoS) parameters (or a combination of them) that are affected by the overall load at the public switch. Additionally, each ISP has to pay a price of $p(y)$ to the IXP per unit traffic, for the use of the public switch. We assume that $d(y)$ and $p(y)$ are given functions (i.e., not part of the strategy); however, we will explore the efficiency of the equilibrium for different forms of the functions $p(y)$ and $d(y)$. For the traffic sent externally, ISP i encounters a per-unit cost of λ_{ij} for traffic exchange with ISP j. This cost may be in terms of additional traffic delays due to longer routes, transit price paid to the ISP's provider, or the cost ISP i incurs for private peering with ISP j. The strategy of each agent (ISP) involves deciding how much of its traffic it should send through the public switch, as opposed to sending externally. In making this decision, we assume that each ISP acts selfishly, focusing on minimizing its own cost. The decisions of ISPs i and j are coupled, and they must agree upon the amount of traffic y_{ij} of the B_{ij} units that is sent through the IXP. Table 1 summarizes some of the most commonly used terms and notations in our model.

2. The private peering can happen at an IXP (if the IXP offers private peering services), or separately.

TABLE 1
Summary of Commonly Used Notation.

Term	Description
y_{ij}	Traffic of ISP pair (i,j) sent publicly through the IXP.
y_i	$\sum_j y_{ij}$, total traffic of ISP i going through the IXP.
y	$\frac{1}{2}\sum_i \sum_j y_{ij}$, total traffic flowing through the IXP.
$\vec{y}$	Total traffic allocation vector (vector of values y_{ij}).
λ_{ij}	Per-unit cost incurred by (i,j) for routing traffic externally.
$d(y)$	Congestion cost per unit traffic incurred at the IX.P
$p(y)$	Price per unit traffic set by the IXP.

Remarks on the model

Note that by 'traffic sent through the IXP', we refer to the traffic sent through the public switch at the IXP. Thus any traffic that is sent through private peering (even if the private peering happens at the same IXP under consideration) is considered a part of the externally routed traffic, i.e., included in $B_{ij} - y_{ij}$ for ISP pair (i,j). Finally, we do not distinguish between the traffic sent from i to j and traffic sent from j to i. In general, an ISP (or the customers of the ISP) benefits from the traffic in both directions, and the two ISPs involved in an exchange must jointly decide whether to exchange this traffic via the IXP or outside of it[3].

In light of the above discussion, we express the price paid by both ISPs i and j, for their traffic exchange at the IXP, as $p(y)y_{ij}$. For easy exposition, in the following, both terms y_{ij} and y_{ji} are utilized, but with the understanding that they represent the same quantity. If the remaining traffic, $B_{ij} - y_{ij}$, is routed through private peering, it is reasonable to assume that the cost of purchasing or leasing any links, ports, etc., to enable this exchange will be proportional to $B_{ij} - y_{ij}$ for both ISPs i and j. Similarly, if this remaining traffic is routed through the ISPs' transit providers, the cost each ISP needs to pay its transit provider, λ_{ij} and λ_{ji}, can be assumed to be proportional to $B_{ij} - y_{ij}$. We first assume that $\lambda_{ij} = \lambda_{ji}$; however, in Section 3.2, we consider a more general model which allows for asymmetric per-unit external routing costs; thus λ_{ij} can be different from λ_{ji}. This allows for possible differences between the two ISPs' transit costs, or their individual costs to privately peer with each other. More generally, asymmetries in the benefits derived from the connection by the two peering ISPs can be resolved with paid peering (as discussed in Section 3.2).[4]

Before detailing our model, we first discuss its relevance to current peering practices. Typically in practice, an ISP pair (i,j) will either send all of their traffic through peering at an IXP, or use the external routing option for all of their mutual traffic. However, our model is more general in that it allows the ISP pair to split their traffic between the two options. This relaxation eases our mathematical discourse and enables us to explore regimes beyond the current practice in traffic exchange between ISPs. Interestingly, from our model it turns out that in the equilibrium solution almost all ISP pairs use only one of the two options (the IXP or external routing), verifying the current practice. In our model, only the IXP pairs whose per-unit external routing cost exactly matches the per-unit effective cost at the switch, end up having to split their traffic between the two options. For large IXPs with lots of ISPs – such that the traffic between any pair of ISPs is very small compared to the overall traffic through the IXP (i.e. $y_{ij} \ll y$) – our modeling assumption holds quite closely. Secondly, while we refer to y_{ij} as the part of the traffic for ISP pair (i,j) that is sent through the IXP, it can also be interpreted as the *equivalent amount of port capacity* that needs to be reserved by both ISPs to carry this traffic through public peering at the IXP. In practice, an ISP i decides (through pairwise agreements) which other ISPs it will peer within an IXP. Accordingly, each ISP individually reserves its port capacity at the IXP so that the port is large enough to carry the traffic to and from all other ISPs it is peering with. There is therefore a direct relationship between the total traffic that an ISP sends through an IXP and the port capacity it reserves (the former being a fraction, say 70% or 80%, of the latter), which allows us to interpret y_i as either of these two terms. The port capacity constitutes the main factor based on which the ISP payments are determined [4], [24]; therefore, either of these interpretations work fine as far as the pricing policy is concerned. Furthermore, even if y_{ij} is interpreted as the traffic between ISP pair (i,j), it should be measured over 'long' time scales. In current practice, the decision on whether or not to publicly peer with another ISP at an IXP is made infrequently – in the timescale of months to years. This implies that when mapping our model to current practice, the traffic rates y_{ij} should be aggregated (averaged) over such time-scales. Furthermore, the congestion (delay) cost $d(y)$ should also be computed (averaged) over such long time-scales, which also represent the timescales at which the pricing strategies (determined by the IXP) and the peering strategies (determined by each pair of ISPs) would be made. Finally, note that peering agreements can depend on additional restrictions, such as traffic volume, traffic ratio, prior customer-provider relationship [47] etc. Such restrictions can be incorporated in our model by removing the specific i,j pairs from consideration that do not satisfy the requirement(s), i.e., setting the corresponding $B_{ij} = 0$.

Some Definitions

Given the above model setup, we next define the *Social Cost* (SC) and *Profit of the IXP* (PX) in order to gain insight into pricing efficiency of the IXP. Overall, SC can be split into the costs incurred by the ISPs at the IXP, and the total payments received by the IXP. The cost for an ISP i, denoted by $C_i(\vec{y}, p(y), d(y))$, is calculated as

$$p(y)\sum_j y_{ij} + d(y)\sum_j y_{ij} + \sum_j (B_{ij} - y_{ij})\lambda_{ij}, \tag{1}$$

where the first and second terms are the costs of sending peering traffic through the public switch – the first is the amount paid to the IXP, and the second is the (implicit) loss of the ISP's revenue caused by the congestion at the switch. The third term is the cost of sending the remaining traffic externally. In Equation 1 we have assumed that the price paid and the congestion cost experienced by each ISP are linear proportional to the corresponding ISP's traffic rate.

3. Although not necessary for the analysis, for the sake of concreteness, y_{ij} can be considered to be either the max or the sum (possibly weighted) of the traffic in the two directions.

4. While we usually assume a fixed per-unit cost λ_{ij}, all our findings also extend to λ_{ij} being a distribution, with different traffic units between i and j having different costs. We omit this model generalization for ease of exposition.

Denoting $c(y) = p(y) + d(y)$ and $L_i(\vec{y}) = \sum_j \lambda_{ij}(B_{ij} - y_{ij})$, the cost of ISP i is

$$C_i(\vec{y}, c(y)) \quad = \quad c(y)y_i + L_i(\vec{y}). \tag{2}$$

Note that $c(y)$ can be viewed as the aggregate cost seen by the ISPs per unit traffic, and therefore equals the sum of the per-unit price charged by the IXP ($p(y)$) and the congestion (delay) cost ($d(y)$). The total cost for all the ISPs is just the summation of C_i for all i. If we denote $\sum_i L_i(\vec{y}) = 2L(\vec{y})$, then the total cost of ISPs becomes

$$C(\vec{y}, c(y)) \quad = \quad 2(c(y)y + L(\vec{y})), \tag{3}$$

where the multiplier of 2 comes from the fact that y_i and y_j both include y_{ij}, i.e., y_{ij} is counted twice.

The IXP receives ISP payments $p(y)$ per unit traffic, and assuming there is no other cost for the IXP, the profit of the IXP equals the revenue earned from the ISP payments, and is expressed as[5]

$$PX(\vec{y}, p(y)) = p(y) \sum_i \sum_j y_{ij} = 2p(y)y. \tag{4}$$

Note that raising the price $p(y)$ does not necessarily increase IXP profit, as higher prices may also result in lower y, as some of the traffic may take the cheaper external routing option.

To define the Social Cost (SC) for our model, we first need to define the cost incurred by the IXP. Since the IXP receives payments from ISPs and has no other cost or revenue, we define the IXP cost as the negative of the IXP profit defined in Equation 4. The social cost of a system is then defined as the sum of the costs incurred by all entities in the system. Thus, we define the social cost (SC) for the given network model as the total cost of the ISPs and IXP cumulatively, which is,

$$\begin{aligned} SC(\vec{y}, d(y)) &= C(\vec{y}, c(y)) - PX(\vec{y}, p(y)), \\ &= 2d(y)y + 2L(\vec{y}) = 2E(y) + 2L(\vec{y}), \end{aligned} \tag{5}$$

where $E(y) = d(y)y$. The first term of this SC is the cost of the congestion at the shared switch in the IXP, and the second is the cost of sending the traffic via external means. Intuitively, both of these components are detrimental to the efficiency of the IXP, and should be minimized. For the rest of the paper, unless otherwise stated, we will assume $E(y)$ to be a continuous, piece-wise differentiable function with $E(0) = 0$ and $E'(y)$ to be a non-decreasing function with $E'(0) = 0$. Note that SC does not consist of $p(y)$ which is the price of per-unit traffic charged by the IXP to the ISPs. However, any change of $p(y)$ will in general affect the traffic flows through the IXP, thereby changing $SC(\vec{y}, d(y))$.

In general, the price $p(y)$ that minimizes SC may not maximize PX (and can in fact result in very low value of PX compared to the optimum profit), and vice versa. It is important to note that *Social Welfare* (SW), which is the weighted sum of utilities across all the entities, has an inverse relationship with social cost. That is to say, if the social cost of a system increases, the social welfare of the system decreases such that the sum of these two (SC + SW) remains constant.

5. Our PoA results on social cost hold even if there is some internal operational cost $r(y)$ which the IXP has, and passes it on to its ISPs by charging each ISP i a value $r(y)y_i/y$. To obtain the same results, we redefine $d(y)$ to be the total of the congestion (delay) cost to the ISPs and the price they are paying to the IXP, and $p(y)$ to be the additional profit that the IXP demands from the ISPs in addition to recovering its operational cost. In other words, we redefine $d(y)$ to be $d(y) + r(y)/y$, and then all the same results hold.

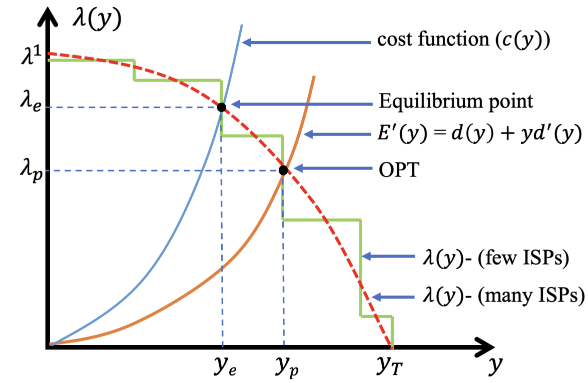

Fig. 1. $\lambda(y)$ curve with $E'(y)$ and $c(y)$ curves.

2.2 Equilibrium Properties

In this paper we assume *price-taking* ISPs, i.e., they see the current cost per unit traffic ($c(y)$), and will only send traffic which is worth paying that cost. From our study of large IXPs, we observed that hundreds of ISPs peer on those (large) IXPs and none of the ISPs are particularly dominant in terms of port capacity [48]. Therefore it is reasonable to assume that each ISP individually does not have enough market power to impact the pricing policy at the IXP. Thus, ISP i will send all the traffic through the IXP as long as $\lambda_{ij} \geq c(y)$, and will route the rest of its traffic externally. Of course, sending more traffic changes the "price" that the ISPs see per unit of traffic, $c(y)$, since it changes y. This leads to the notion of of equilibrium traffic flow, defined as follows:

Definition 2.1. *A traffic flow $\vec{y}_e$ with $y_e = |\vec{y}_e|$ is said to be an equilibrium flow if and only if all the traffic with $\lambda_{ij} > c(y_e)$ is sent through the IXP and the traffic with $\lambda_{ij} < c(y_e)$ is not.*

Based on this definition, we simplify our terminology and use the term 'equilibrium traffic flow' to refer to the *total* flow through the IXP at equilibrium (a scalar), and denote it by y_e.

Next, we state two important properties of equilibrium traffic flows that will be useful in our PoA analysis. We first define the notion of the *inverse demand curve*, denoted as $\lambda(y)$. Illustrated by Figure 1, this curve is constructed as follows. First, the λ_{ij} values are arranged in a decreasing order (ties broken arbitrarily); let λ^k be the k^{th} highest value, and B^k be the corresponding traffic demand. Then, the $\lambda(y)$ curve is a non-increasing step-function, with the k^{th} step having height of λ^k and a width of B^k. Let $\lambda(y^-)$ denote the limit of $\lambda(x)$ as x approaches y from below, and similarly $\lambda(y^+)$ if it approaches y from above. Then $\lambda(y)$ can be interpreted as the cost of sending traffic through the IXP (as seen by an ISP) when the total IXP traffic is y. Therefore, as $\lambda(y)$ increases, y decreases. Since a traffic is sent through the IXP only if its external routing cost is greater than the cost of sending through the IXP, y corresponds to the sum of all ISP traffic whose external routing cost is greater than $\lambda(y)$. We have the following property:

Theorem 2.1. *y_e is an equilibrium traffic flow if and only if $\lambda(y_e^-) \geq c(y_e) \geq \lambda(y_e^+)$. Moreover, such a flow always exists.*

Proof Outline of Theorem 2.1: This follows directly from definition of equilibrium and the $\lambda(y)$ curve: equilibrium traffic sends demand with highest λ_{ij} first, until it equals the cost $c(y)$. Due to the properties of $\lambda(y)$, and $c(y) > 0$ for $y > 0$, at least one such intersection point exists. □

Notice also that multiple equilibria may exist. First, there could be several flow amounts y_e with $\lambda(y_e^-) \geq c(y_e) \geq \lambda(y_e^+)$ if the function $c(y)$ is not strictly increasing. Second, even for a fixed total flow y_e, there could be several different traffic pairs with equal $\lambda_{ij} = c(y_e)$ values, and sending any subset of them (as long as the total flow amount equals y_e) will yield an equilibrium flow.

Finally, we derive an important property of the optimal traffic vector (OPT), one that minimizes the total social cost, with which the equilibrium solution will be compared. Again, with slight abuse of terminology, by 'optimal traffic flow' we refer to the *total* traffic flow at the social optimum, denoted by y_p; and in such a traffic vector the traffic with largest λ_{ij} will be sent in order to minimize social cost.

Theorem 2.2. *At social optimality, all the traffic with $\lambda_{ij} > E'(y_p)$ flows through the IXP and all traffic with $\lambda_{ij} < E'(y_p)$ does not. Also, $\lambda(y_p^-) \geq E'(y_p) \geq \lambda(y_p^+)$).*

Proof Outline of Theorem 2.2: To see why this is true, consider a traffic vector $\vec{y_t}$ with $y_t = |\vec{y_t}|$ so that the traffic with largest λ_{ij} is sent. Then, from the definition of social cost, we have that:

$$SC(\vec{y_t}, d(y_t)) = 2\int_0^{y_t} E'(y)dy + 2\int_{y_t}^{y_T} \lambda(y)dy.$$

The above applies for both $y_t = y_e$ and $y_t = y_p$. To minimize social cost, it is clear that y_p should be the intersection of the λ and E' curves (see Figure 1). □

3 Social Cost Analysis

In this section, we analyze the Price of Anarchy (PoA) for the traffic exchange game defined in Section 2, calculated as the ratio of social cost (SC) at the worst equilibrium to the SC at the optimal solution (OPT) that minimizes social cost. We first show that under an "optimal" pricing scheme, the PoA equals unity, i.e., all equilibria of the traffic exchange game attain social optimality (Theorem 3.1). We then analyze the PoA attained by two other natural pricing policies, under two broad classes of delay functions.

Theorem 3.1. *The pricing policy $p(y) = d'(y)y$ attains a PoA of 1 (i.e., minimizes Social Cost).*

Proof: The proof of Theorem 3.1 is very straightforward. For $p(y) = d'(y)y$, we have $c(y) = E'(y)$; therefore the equilibrium y_e and OPT y_p are the same, implying PoA=1. □

However, using the social-cost optimal pricing policy (as given by Theorem 3.1) can result in very poor profit. To see this, consider a simple (linear) external routing cost curve $\lambda(y) = 1 - y$, where $y_T = 1$. Now let $d(y)$ be small enough (which also makes $E'(y)$ very small) such that y_p is 0.99 (refer to Figure 1), and hence $\lambda(y_p) = 0.01$. Then the social-cost optimal policy ($y_e = y_p$) results in an IXP profit of at most 0.0099, since the price the IXP charges is $p(y_p) \leq c(y_p) = 0.01$ and 0.99 units of traffic pay this price. However, the maximum achievable profit is about 0.25, which is attained when the per-unit price is chosen to be about 0.5 resulting in equilibrium traffic of $y_e = 0.5$. So, the ratio between the maximum achievable profit to the achieved profit for this pricing scheme (which will be defined as PoA of profit in later sections) is 25.25. Now, say we set the price instead so that it results in a slightly lower equilibrium traffic $y_e = 0.95$ than in the socially optimum solution; this corresponds to a price that is slightly lower than 0.05. Then we get the PoA of profit that is approximately 5.26, which is obviously a lot better than before, without sacrificing much social welfare. This motivates us to look for alternative pricing policies that may be slightly sub-optimal in terms of social cost, but result in good profit. Our consideration of proportional pricing, which is analyzed next, is guided by this.

3.1 PoA Analysis for Social Cost - Proportional Pricing

Definition 3.1. *Proportional pricing with a proportionality factor $\beta \geq 1$ has a per-unit price $p(y)$ defined as $p(y) = (\beta - 1)d(y)$. In other words, the effective cost seen by the ISPs sharing the IXP per unit traffic is $c(y) = p(y) + d(y) = \beta d(y)$.*

Definition 3.2. *Zero pricing has a pricing function $p(y) = 0$, thus making the effective per-unit cost for ISPs consist only of congestion cost: $c(y) = d(y)$.*

Clearly, zero pricing can be viewed as a special case of proportional pricing with $\beta = 1$.

In the PoA analysis for social cost that follows next, we consider two broad classes of congestion cost (delay) functions: 1) polynomial delay functions, 2) queuing delay functions. Proofs of the main results in this section are provided in Appendix B, unless otherwise mentioned.

3.1.1 PoA for Social Cost under Polynomial Delay Functions

The PoA for social cost in our model crucially depends on the properties and convexity of the congestion cost function $d(y)$. We make no assumptions about the λ_{ij} distribution, but only about the congestion cost functions. We begin by considering congestion cost (delay) functions which exhibit polynomial growth rates.

Theorem 3.2. *For proportional pricing (i.e., $c(y) = \beta d(y)$), if congestion cost (delay) function $d(y) = ay^n$ with $a > 0, n \geq 1$, and*
i) $\beta \leq n+1$, then PoA is bounded by $[\beta - n(\frac{\beta}{n+1})^{(n+1)/n}]^{-1} \leq \frac{n+1}{\beta}$;
ii) $\beta > n+1$, then PoA is bounded by $\frac{\beta}{n+1}[\frac{\beta n}{(\beta-1)(n+1)}]^n \leq \frac{\beta}{n+1}$.

Corollary 3.2.1. *For zero pricing (i.e., $c(y) = d(y)$), if congestion cost (delay) function $d(y) = ay^n$ for some constant $a > 0$, then the PoA is bounded by $(1 - n(1+n)^{-(n+1)/n})^{-1}$.*

Corollary 3.2.2. *If the delay cost, $d(y)$, satisfies $\frac{d}{dy}(by^n) \leq d'(y) \leq \frac{d}{dy}(ay^n)$ for some positive constants a, b; then the PoA* *bounds of Theorem 3.2 and Corollary 3.2.1 hold with an additional multiplicative factor of* $\gamma = \frac{a}{b}$.

Corollary 3.2.1 follows directly from part i) of Theorem 3.2, and taking $\beta = 1$. Corollary 3.2.2 shows how our results generalize when the congestion cost (delay) function can be sandwiched between two polynomial functions with the same exponent n.

We can derive additional bounds on the PoA for social cost under certain stronger assumptions on the congestion cost (delay) function, as stated below.

Lemma 3.1. *For proportional pricing (i.e.,* $c(y) = \beta d(y)$*), if congestion cost (delay) is any convex function with* $d(y) \le ay^n$*, then with* $\beta > n+1$*, the PoA is bounded by* $\frac{\beta}{2}$.

When $\beta \le n+1$, a better bound than Corollary 3.2.2 can be found (proof in Appendix D, included in the Supplementary Material).

Theorem 3.3. *For proportional pricing (i.e.,* $c(y) = \beta d(y)$*), If the congestion cost (delay) function* $d(y)$ *satisfies* $\frac{d}{dy}(by^n) \le d'(y) \le \frac{d}{dy}(ay^n)$ *for any* $a > b > 0$*, then with* $\beta \le n+1$*, the PoA is bounded by* $(\frac{\gamma\rho_m}{(\gamma-1)+\rho_m^{-n}} + \beta(1-\rho_m))^{-1}$*, where* $\rho_m = (\frac{2\beta}{\gamma(n+1)-2\beta(\gamma-1)+\sqrt{\gamma^2(n+1)^2-4\beta\gamma n(\gamma-1)}})^{1/n}$ *and* $\gamma = \frac{a}{b}$.

Characteristics of the PoA bounds: The PoA bounds for different values of n with $d(y) = ay^n$ are shown for zero and proportional pricing in Figure 2. We can see that, although the equations of the bounds were quite complicated, both the bounds are quite well-behaved. With zero pricing (which is a special case of proportional pricing with $\beta = 1$), if the delay cost is a linear function ($n = 1$), then the PoA is 1.33, which means irrespective of the shape of the $\lambda(y)$ curve (i.e., the values of the external routing costs of ISPs and the IXP), the worst equilibrium will only cost 33% more than the optimum cost. The results also show that the social cost benefits of proportional pricing (for a well chosen β value) over zero pricing can be quite significant.

For proportional pricing, the case is a bit more complicated with two variables n and β, still the bounds exhibit a simple linear-like behavior. If the value of n is increased from 1, then with $\beta > n+1$, the PoA starts to decrease and goes to 1 when $\beta = n+1$; after that with $\beta < n+1$ the PoA starts to increase. The value of PoA becoming 1 at $\beta = n+1$ coincides with the cost function $c(y)$ becoming equal to $E'(y)$. Overall, the PoA for social cost remains very small for reasonable values of n and β, showing that it is possible for IXPs to make a nice profit while still attaining a traffic exchange solution between ISPs that is close to optimal in terms of social cost. We explore this further in Section 4 when we analyze IXP profit.

3.1.2 PoA for Social Cost under Queuing Delay Functions

Next we consider the queuing delay function[6] (modeling the public switch at the IXP as a single server) expressed as:

$$d(y) = \frac{1}{\mu - y}, \tag{6}$$

6. Note that the queuing delay function we consider represents that of the basic M/M/1 queue. The PoA results easily generalize to functions up to a multiplicative constant, i.e., delay functions of the form $a/(\mu - y)$; extension of the results to other more general forms of delay functions remains open for future investigation.

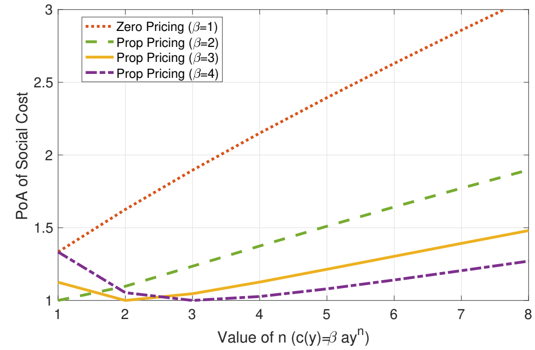

Fig. 2. PoA bounds for SC (polynomial delay function).

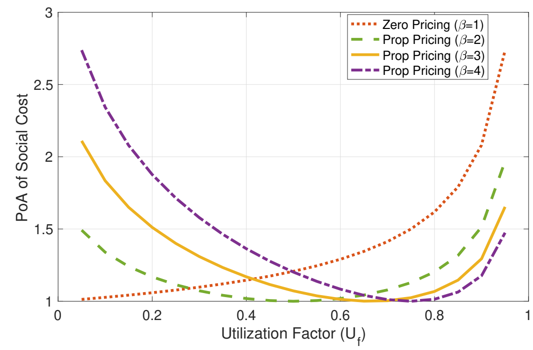

Fig. 3. PoA bounds for SC (queuing delay function).

where μ is the processing rate of traffic at the public switch. Note that $d(y)$ can become unbounded as the aggregate traffic load on the switch approaches capacity μ. These type of congestion cost (delay) functions are not analyzable in the same framework as polynomial delay functions considered earlier; they need to be treated separately, as we do here. Normally, switches are operated under 60-70% of the full capacity, as otherwise congestion delays would become too high to exchange traffic smoothly. This fact will be utilized in deriving the PoA for social cost.

Definition 3.3. *Utilization factor,* U_f*, satisfying* $0 < U_f < 1$*, is the ratio of traffic load on a network to the total capacity of the network. If the total traffic at equilibrium is* y_e*, then we define it to be* $U_f = y_e/\mu$.

The following results provide the PoA bounds for social cost under proportional pricing and zero pricing with queuing delay functions. The result for zero pricing again follows as a special case ($\beta = 1$) of the proportional pricing result.

Theorem 3.4. *For proportional pricing (i.e.,* $c(y) = \beta d(y)$*) and congestion cost (delay) function* $d(y) = \frac{1}{\mu - y}$*, the PoA is bounded by*

i) $\frac{U_f\sqrt{\frac{1-U_f}{\beta}}}{(1-U_f)[2-\sqrt{1-U_f}(\frac{1+\beta}{\sqrt{\beta}})]}$*, when* $U_f \ge 1 - \frac{1}{\beta}$*;*

ii) $\frac{\left(\sqrt{\beta}-\sqrt{U_f(\beta-1)}\right)^2}{1-U_f}$*, when* $U_f < 1 - \frac{1}{\beta}$.

Corollary 3.4.1. *For zero pricing (i.e.,* $c(y) = d(y)$*) and congestion cost (delay) function* $d(y) = \frac{1}{\mu - y}$*, the PoA is bounded by* $\frac{U_f\sqrt{1-U_f}}{2(1-U_f)(1-\sqrt{1-U_f})}$.

Characteristics of the PoA bounds: Figure 3 shows the PoA upper bound with the increase of utilization factor, U_f of the switch. We see that the upper bound on PoA becomes 1 when $U_f = 1 - 1/\beta$. Also, PoA upper bound increases at an exponential rate on both sides of $U_f = 1 - 1/\beta$, but the rate of increment is higher on the side where $U_f > 1 - 1/\beta$. Typically, routers and switches maintain a utilization factor in the range of 40% to 70%, and for that operational range we see that PoA is below 1.2 for $\beta = 2$ and 3. This means that the worst equilibrium has social cost only 20% higher than the optimum under normal operating loads at the IXP.

3.2 PoA under Asymmetric External Routing Costs

All the results given until now had the assumption of $\lambda_{ij} = \lambda_{ji}$. However, as mentioned earlier, it is possible that the two ISPs i and j may encounter different per-unit costs for routing traffic externally between themselves, due to differences in transit pricing, or cabling/leasing cost associated with private peering. In this subsection, we will discuss the effect of $\lambda_{ij} \neq \lambda_{ji}$ on the PoA bounds calculated. The consideration of asymmetry between the external routing costs makes the proofs substantially longer and more difficult, and are provided in Appendix E of the Supplementary Material.

The consideration of asymmetric external routing costs requires us to revisit the definition of equilibrium. From the discussion of Section 2.1, we know that when an ISP i has some traffic to exchange with ISP j, it compares the value λ_{ij} of that traffic with the cost $c(y)$ of exchanging that traffic via the IXP. As for the current case $\lambda_{ij} \neq \lambda_{ji}$, suppose that $\lambda_{ij} > c(y) > \lambda_{ji}$. Hence, ISP i will want to exchange traffic (since, then, its cost will decrease by λ_{ij} and increase by $c(y)$), whereas ISP j will not want to (since, then, its cost will decrease by λ_{ji} and increase by $c(y)$). Since both participants are needed to form a peering connection at an IXP, the traffic exchange will not happen in this case. Thus, in an equilibrium state, exchange of traffic between two ISPs is only possible when $\min(\lambda_{ij}, \lambda_{ji}) > c(y)$.

Definition 3.4. *A traffic flow $\vec{y}_n$ with $y_n = |\vec{y}_n|$ is said to be an equilibrium flow for the case of $\lambda_{ij} \neq \lambda_{ji}$, when all the traffic with $\min(\lambda_{ij}, \lambda_{ji}) > c(y_n)$ is sent and the traffic with $\min(\lambda_{ij}, \lambda_{ji},) < c(y_n)$ is not sent.*

We will now jump directly to the PoA results for the three cases: socially optimal pricing (Theorem 3.5), and proportional pricing (Theorem 3.6 and Theorem 3.7, which discusses the cases of polynomial and queuing delay functions, respectively). Recall that $c(y) = p(y) + d(y)$, and note that PoA results for zero pricing follows from Theorems 3.6 and 3.7 by setting $\beta = 1$.

Theorem 3.5. *If $c(y) = E'(y)$, i.e., $p(y) = d'(y)y$, when $\lambda_{ij} \neq \lambda_{ji}$ with $\alpha = \max\{\frac{\lambda_{ij}}{\lambda_{ji}}\}$, PoA is bounded by $\frac{1+\alpha}{2}$.*

Theorem 3.6. *If $c(y) = \beta d(y)$ (i.e., proportional pricing) and $d(y) = ay^n$, when $\lambda_{ij} \neq \lambda_{ji}$ with $\alpha = \max\{\frac{\lambda_{ij}}{\lambda_{ji}}\}$, PoA is bounded by*
I) $(\frac{1+\alpha}{2})(\frac{\beta}{n+1}) \times$ (PoA bounds for $\lambda_{ij} = \lambda_{ji}$); when $\beta > n+1$;
II) $(\frac{1+\alpha}{2}) \times$ (PoA bounds for $\lambda_{ij} = \lambda_{ji}$); when $\beta \leq n+1$.

Theorem 3.7. *If $c(y) = \beta d(y)$ (i.e., proportional pricing) and $d(y) = \frac{1}{\mu - y}$, when $\lambda_{ij} \neq \lambda_{ji}$ with $\alpha = \max\{\frac{\lambda_{ij}}{\lambda_{ji}}\}$, PoA is bounded by*
I) $\frac{1+\alpha}{2} \times \beta(1 - U_f) \times$ (the PoA bound when $\lambda_{ij} = \lambda_{ji}$) ; when $U_f < (1 - \frac{1}{\beta})$;
II) $\frac{1+\alpha}{2} \times$ (the PoA bound when $\lambda_{ij} = \lambda_{ji}$); when $U_f \geq (1 - \frac{1}{\beta})$.

Roughly speaking, Theorems 3.5-3.7 state that when the notion of equilibrium is defined as in Definition 3.4, the PoA bounds for the asymmetric external routing cost scenario differs from the symmetric case by about $\frac{1+\alpha}{2}$, where α represents the degree of asymmetry between the costs. Note that in the above theorems, the degree of asymmetry α is calculated only over ISP pairs i, j at the IXP that are interested in sending traffic to one another, i.e., $B_{ij} > 0$.

Paid Peering. Note that the PoA bounds in Theorems 3.5-3.7 are large when the degree of asymmetry in the external routing costs is high. However, our simulation results in Section 5 show that, in practice, the efficiency at equilibrium compared with the optimum is typically much better than these worst case bounds. Further, when one ISP encounters a much higher cost than the other in enabling traffic exchange between the two, *paid peering* would make sense. Indeed, paid peering has been suggested by some as a solution to growing asymmetry in costs experienced (or benefits realized) between two peering ISPs [15]. The following result states that in the case of asymmetric costs, if every pair of ISPs share the external routing costs via *Nash Bargaining* [49], it results in the same PoA as in the symmetric case. This is because if i pays λ_{ij} and j pays λ_{ji} (where $\lambda_{ji} > \lambda_{ij}$, say) in external routing costs, it is not difficult to see that the Nash Bargaining solution would result in a payment ("settlement") of $(\lambda_{ji} - \lambda_{ij})/2$ from i to j. Thus, when using a Nash Bargaining protocol, one ISP would pay the other exactly enough so that the effective external costs for not using the IXP would become equal, thus resulting in the same bounds as if λ_{ij} were always equal to λ_{ji}.

Theorem 3.8. *If ISPs are allowed to pay each other (paid peering), and use Nash Bargaining to determine payments, then the PoA bounds are the same as in the symmetric case (i.e., as in Section 3.1).*

The above discussion suggests a specific paid peering policy (based on Nash Bargaining) that results in efficient traffic equilibria. However, broader arguments on paid peering to recover traffic sensitive costs such as those described in [14] are outside the scope of this work.

3.3 PoA Analysis for Social Cost - Constant Pricing

We conclude this section on social cost analysis by analyzing the PoA under constant per-unit pricing, which will be contrasted with proportional pricing that we have analyzed so far. Proofs are provided in Appendix F of the Supplementary Material.

Definition 3.5. *Constant pricing is the pricing policy where the per unit price is a constant value (γ). Thus, we have $p(y) = \gamma$ and $c(y) = \gamma + d(y)$.*

In most of our analysis with constant pricing, we will assume that $\gamma = r \times \lambda^1$, where $0 < r < 1$, and λ^1 is the largest external routing cost among the ISP pairs. For constant pricing, the PoA of social cost ($PoA(SC)$), depends on the traffic flowing through the IXP. Defining y_q to be the traffic for which $c(y_q) = E'(y_q)$, and recalling that y_p represents the socially optimal traffic flow (see Theorem 2.2), we obtain the following bounds:

Theorem 3.9. *For constant Pricing (i.e., $c(y) = \gamma + d(y)$), if congestion cost (delay) function $d(y) = ay^n$ with $a > 0, n \geq 1$, and*
i) $y_e \leq y_p \leq y_q$, then $PoA(SC)$ is bounded by $[\rho^{n+1} + (1-\rho)(\frac{\gamma}{ay_p^{n+1}} + \rho^n)]$, where $\rho = y_e/y_p \leq 1$;
ii) $y_e > y_p > y_q$, then $PoA(SC)$ is bounded by $[\rho^{n+1} + (1-\rho)(\frac{\gamma}{ay_e^{n+1}} + \rho^n)]^{-1}$, where $\rho = y_p/y_e \leq 1$.

From Theorem 3.9, it can be observed that the $PoA(SC)$ depends on the traffic values (i.e., y_e, y_p) and thus depends on the $\lambda(y)$ curve. So, a proper bound on $PoA(SC)$ for constant pricing with polynomial delays cannot be determined without having further knowledge about external routing costs (λ values) of the participating ISPs. The information on these λ values are private to the ISPs, and therefore not generally known to the IXP. Even if the IXP is interested in choosing the price γ to (approximately) minimize the social cost, this poses a practical difficulty in implementing this constant price that achieves a small $PoA(SC)$.

Theorem 3.10. *For constant Pricing (i.e., $c(y) = \gamma + d(y)$), if congestion cost (delay) function $d(y) = \frac{1}{\mu - y}$, and*
i) $y_e \leq y_p \leq y_q$, the PoA is bounded by $\frac{\rho(1-\rho)+\gamma\mu(\rho-U_f)(1-\rho)(1-U_f)}{\rho(1-U_f)}$, where $\rho = \sqrt{\frac{\gamma\mu U_f(1-U_f)}{1+\gamma\mu(1-U_f)}}$;
ii) $y_e > y_p > y_q$, the PoA is bounded by $\frac{U_f(1-\rho)}{\rho(1-U_f)+(U_f-\rho)(1-\rho)[1+\gamma\mu(1-U_f)]}$, where $\rho = 1 - \sqrt{\frac{1-U_f}{1+\gamma\mu(1-U_f)}}$.

Theorem 3.10 represents a better bound than Theorem 3.9 in terms of dependency. Unlike Theorem 3.9, Theorem 3.10 depends on utilization factor U_f and the constant per-unit price γ. Since the range of γ ($< \lambda^1$) and U_f may be known or estimated, then we can derive a range of $PoA(SC)$ values from Theorem 3.10. Even then, with constant pricing, a small $PoA(SC)$ over a wide range of $\lambda(y)$ curves (generally unknown to the IXP) seems difficult to attain, as we will see in Section 5.

4 IXP Profit Analysis

Towards analyzing IXP profit at equilibrium, we define the Price of Anarchy for profit in the same way as we defined the PoA of social cost. Proofs of the results for the proportional pricing case (Section 4.1) are provided in Appendix C, while those for the constant pricing case (Section 4.2) are in Appendix G of the Supplementary Material.

Definition 4.1. *The Price of Anarchy for profit ($PoA(PX)$) is the ratio of maximum achievable profit to the profit achieved at equilibrium for some pricing scheme.*

Similar to $PoA(SC)$, note that $PoA(PX)$ can only have values greater than or equal to 1.

4.1 PoA Analysis for Profit - Proportional Pricing

To obtain bounds on PoA for profit, we define a parameter K such that $\lambda^1 \leq K\lambda(y_e) = K\lambda_e$. This implies that all the traffic having external routing cost values greater than $\frac{\lambda^1}{K}$ are exchanged through the IXP. For example, with K=4, all traffic having an external routing cost more than 25% of λ^1 (the maximum external routing cost) are exchanged through the IXP.

Theorem 4.1. *For proportional pricing (i.e., $c(y) = \beta d(y)$), if congestion cost (delay) function $d(y) = ay^n$ with $n \geq 1$, and*
i) $\beta \leq n+1$, then PoA is bounded by $\frac{K\beta\rho - \rho^{n+1}}{\beta - 1}$, where $\rho = min((\frac{K\beta}{n+1})^{1/n}, 1)$;
ii) $\beta > n+1$, then PoA is bounded by $\frac{\beta(K-1) + \frac{\beta n}{n+1}(\frac{\beta}{n+1})^{1/n}}{\beta - 1}$.

Characteristics of the PoA bounds: Figure 4 shows the PoA bounds for IXP profit under polynomial delay functions, as given by Theorem 4.1. Since IXP profit is directly related to the value of β multiplied by $d(y)$, we plot the PoA results with respect to different values of β. The first noticeable property of $PoA(PX)$ is that it depends on (increases with) the value of K (where $\lambda^1 \leq K\lambda_e$), as expected from Theorem 4.1. We also observe that the $PoA(PX)$ values tend to decrease with increase in n. Lastly, with increase of β we see a decreasing trend in $PoA(PX)$ for smaller values of β, which is expected since profit at equilibrium is proportional to β. However, as β is increased further beyond a point, the equilibrium traffic (y_e) decreases sharply to compensate for this effect, resulting in an overall reduction in $\beta \times d(y)$, and thereby a decrease in PX (increase in $PoA(PX)$). From the plots, we see that a range of $\beta = 3$ to 5 attains near-optimal $PoA(PX)$ values.

Theorem 4.2. *For proportional pricing (i.e., $c(y) = \beta d(y)$), if congestion cost (delay) function $d(y) = \frac{1}{\mu - y}$, and*
i) $U_f \leq 1 - \frac{1}{\beta}$, then PoA is bounded by $\frac{\beta U_f(K-1)+(\sqrt{1-U_f}-\sqrt{\beta})^2}{U_f(\beta-1)}$;
ii) $U_f > 1 - \frac{1}{\beta}$, then PoA is bounded by $\frac{(\sqrt{K\beta}-\sqrt{1-U_f})^2}{U_f(\beta-1)}$.
(In both cases it was assumed that $\lambda^1 \leq K\lambda_e$.)

Characteristics of the PoA bounds: Figure 5 shows the PoA bounds for IXP profit under queuing delay functions, as given by Theorem 4.2. Just like Figure 4, all the $PoA(PX)$ values are plotted against the values of β. The first noticeable property for the PoA bound is the same as for polynomial delay, that is, PoA depends heavily on the value of K (where $\lambda^1 \leq K\lambda_e$). Then we see an almost inverse relationship trend of $PoA(PX)$ with utilization factor (U_f); the more the value of the utilization factor is, the lower the $PoA(PX)$. As a final observation we notice that the effect of β on the $PoA(PX)$ is quite similar to the case of polynomial delay, however with the queuing delay functions, the range of $beta$ which gives near-optimal $PoA(PX)$ is much wider (from $\beta = 3$ to 10 and even beyond).

4.2 PoA Analysis for Profit - Constant Pricing

Theorem 4.3. *For constant pricing (i.e., $c(y) = \gamma + d(y)$), if congestion cost (delay) function $d(y) = ay^n$ with $a > 0, n \geq 1$, and*

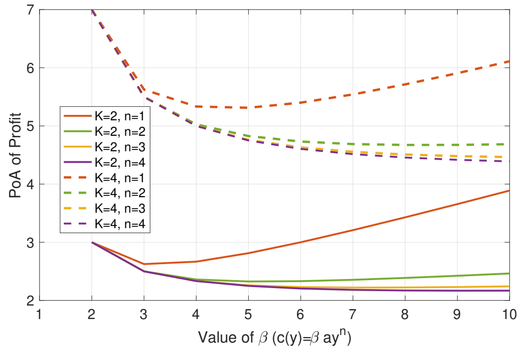

Fig. 4. PoA bounds for profit (polynomial delay func.).

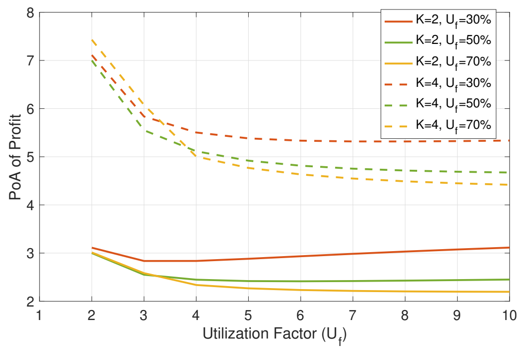

Fig. 5. PoA bounds for profit (queuing delay function).

i) $y_e \le y_p \le y_q$, *then PoA is bounded by* $(Ky_e + y_p - y_e)(\frac{1}{y_e} + \frac{ay_e^{n-1}}{\gamma}) - \frac{ay_p^{n+1}}{\gamma y_e}$;
ii) $y_e > y_p > y_q$, *then PoA is bounded by* $\frac{\rho}{\gamma}[K\gamma + ay_p^n(K-1)]$, *where* $\rho = y_p/y_e \le 1$.

Similar to Theorem 3.9, we observe that $PoA(PX)$ under constant pricing with polynomial delay function also depends heavily on the traffic values (i.e., y_e, y_p), which in turn depends on the $\lambda(y)$ curve. So unlike proportional pricing, providing a good bound on $PoA(PX)$ for constant pricing with polynomial delay function requires further knowledge about external routing costs (λ values) of the participating ISPs.

Theorem 4.4. *For constant pricing (i.e.,* $c(y) = \gamma + d(y)$*), if congestion cost (delay) function* $d(y) = \frac{1}{\mu - y}$*, and*
i) $y_e \le y_p \le y_q$, *then PoA is bounded by* $\frac{(1-\rho)[(K-1)U_f+\rho][\gamma\mu(1-U_f)+1]-\rho(1-U_f)}{\gamma\mu U_f(1-U_f)(1-\rho)}$,
where $\rho = 1 - \sqrt{\frac{1-U_f}{1+\gamma\mu(1-U_f)}}$;
ii) $y_e > y_p > y_q$, *then PoA is bounded by* $\frac{K\rho(1-\rho)[\gamma\mu(1-U_f)+1]-\rho(1-U_f)}{\gamma\mu U_f(1-U_f)(1-\rho)}$,
where $\rho = 1 - \sqrt{\frac{1-U_f}{K[1+\gamma\mu(1-U_f)]}}$.

Observations similar to those in the $PoA(SC)$ case (see Theorem 3.10) can be made here as well, i.e., if we know the range of γ and U_f, then we can derive a range of $PoA(PX)$ values from Theorem 4.4.

5 Simulation Results

5.1 Data Collection

To achieve realistic traffic demand values B_{ij} and external routing costs λ_{ij}, data from PeeringDB and CAIDA databases were collected and analyzed. We used PeeringDB to get information about the locations of the IXPs, the ISPs peering in that location (also called Point-of-Presence (PoP)) and the port capacity each ISP has purchased. A short summary of the current statistics of ISPs and IXPs in the USA is given in Table 2; note that an IXP can constitute of multiple facilities, which are typically located close to one another. Moreover, a map of the USA showing the location of the IXPs and the number of ISPs in those IXPs are shown in Figure 7. Additionally, it is noteworthy that PeeringDB provides comprehensive details on a given ISP, including its name, organizational structure, type, and peering policy, among several other fields of information. According to PeeringDB there are currently 745 access ISPs, 473 content ISPs, and 686 transit ISPs active in the US mainland. All the data available in PeeringDB is stored in the JSON format and undergoes regular updates.

On the other hand, we utilized CAIDA to get the number of active routers and their approximate location (at a city level) for each ISP, to approximate the amount of traffic that may be generated for that ISP at that location. We downloaded these data from the Internet Topology Data Kit (ITDK) database of CAIDA. The nodes.as text file contained the ASN to router ID mapping, while the nodes.geo text file contained the location information of the router IDs. Furthermore, we leveraged the CAIDA rank database to classify ISPs into tiers.

5.2 Simulation Setup

5.2.1 Generating External Routing Cost (λ) values

To generate the $\lambda(y)$ curves we need two sets of values: i) the traffic demand between ISPs (B_{ij}), and ii) the per-unit external routing costs (λ_{ij}) for that traffic. While the exact values of these are very difficult to estimate closely, we make several reasonable approximations based on the PoP locations (obtained from PeeringDB), router densities (obtained from CAIDA), and previously published models on traffic demand and pricing. The traffic demand between two ISPs serving at two different PoP locations is determined using the gravity model [50]. If ISP i has R_A number of routers serving at location A and ISP j has R_B number of routers serving at location B, then the traffic demand between these two ISPs for these two locations is thus approximated as $Y_{AB} = \frac{R_A \times R_B}{d_{AB}^2}$. Then summing up these traffics for all possible pair of locations gives us the total traffic demand between two ISPs, hence $B_{ij} = \sum_{A,B} Y_{AB}$. To calculate the external routing cost for B_{ij} we follow [44], which models transit costs as being linearly or logarithmically proportional to the distance that traffic has to travel. Since traffic between different locations of the same ISP pair (say Y_{AB}) is going to travel different distances (d_{AB}), we use the weighted average of these distances: for some ISP pair (i,j), we set $d_{ij} = \frac{\sum_{A,B} Y_{AB} d_{AB}}{\sum_{A,B} Y_{AB}}$. Thus, we have the per-unit external routing cost as, $\lambda_{ij} = \Delta \times d_{ij}$ or $\lambda_{ij} = \Delta \times log(d_{ij})$, for an appropriately chosen constant Δ.

Note that if ISP pair (i,j) decides to peer, then the B_{ij} traffic may be split across the different PoP locations that the two ISPs have in common. To find how much of B_{ij} traffic of the ISP pair (i,j) will flow through each of these common points, we used three different approaches.

TABLE 2
Summary on IXPs and ISPs in the USA

Total IXPs (PoP locations)	140
Total Facilities	1078
Facilities supporting public peering only	293
Facilities supporting both type of peering	830
Total ISPs	2821
ISPs (public peering only)	723
ISPs (private peering only)	728
ISPs (both types of peering)	1370

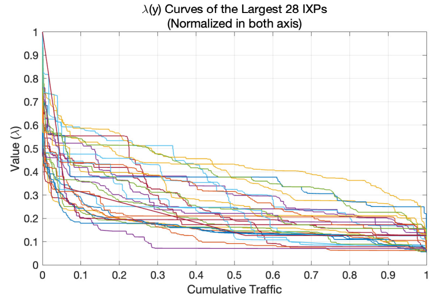

Fig. 6. Generated $\lambda(y)$ curves.

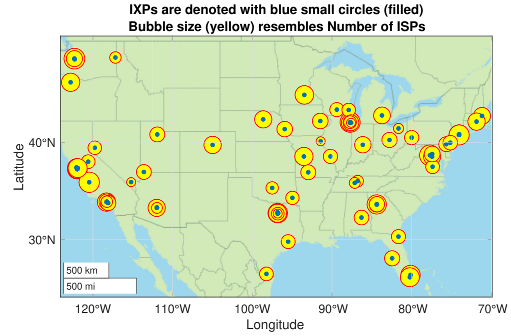

Fig. 7. USA map with IXP locations.

The first approach sends all the traffic through the IXP that minimizes the total end-to-end geographical distance; the second approach divides the traffic equally among all common IXP locations; and third approach splits the traffic as inversely proportional to total end-to-end geographical distance of each path. Although all these methods yielded different traffic values at the IXPs, the nature of the external routing cost curves ($\lambda(y)$ curves) were quite limited in all these three cases. Since our performance results mainly depend on the shape of these $\lambda(y)$ curves and they ended up being similar; therefore, in the following we only present the results for the third approach. Figure 6 shows $\lambda(y)$ curves generated with the third approach for the largest (in terms of number of participating ISPs) 28 IXPs in USA.

5.2.2 *PoA Simulation*

To obtain the PoA values for social cost and profit, simulations were done for the largest 28 IXPs among the 140 IXPs present in USA. Most of the remaining (smaller) IXPs have a very small number of participating ISPs, resulting in a few discrete $\lambda(y)$ values and making the study of the equilibrium uninteresting. Also, from the PeeringDB port capacity data, it was found that more than 95% of the total port capacities (which can be seen as an indicator of the traffic flowing through these IXPs) are accounted for by considering the largest 28 IXPs. To find the PoA values for proportional and constant pricing with polynomial delay functions, we used different values of a, where $d(y) = ay^n$, and then the PoA value was calculated considering the delay function (or, equivalently, the value of a) that resulted in the worst PoA. Since the PoA value also critically depends on the $\lambda(y)$ curves which will differ from IXP to IXP, both worst and average case PoA values were calculated by taking the the worst value and average values over all the $\lambda(y)$ curves, respectively. For the case of queuing delay functions, the $\lambda(y)$ curves were normalized with respect to the value of μ (recall that $d(y) = \frac{a}{\mu - y}$) and then results for different utilization factors (U_f) were generated.

5.3 Results and Discussion

5.3.1 *Comparing with Theoretical bounds*

The maximum value of PoA (which we denote as $Max(PoA)$) for proportional pricing obtained from simulation are plotted against their corresponding theoretical bounds in Figures 8, 9, 11, and 12. In all these Figures, the curves marked $Theo$ represents the corresponding theoretical bounds; whereas the curves marked Sim are the PoA values found through simulation. Also, most of the profit bounds and simulation results presented in this section are for $K = 2$. Based on our analysis of the inverse demand curves from Figure 6, we observed that $K = 2(\lambda^1 \leq K\lambda_e$, from definition of K in Section 4.1) results in around 30-40% of the total traffic data flowing through the IXP in equilibrium. This finding aligns with previous studies, such as [13], which reported that 20% of global internet traffic traversed through IXPs in 2016, and extrapolating their data leads us to a value of approximately 30% for 2022.

We note that the $Max(PoA)$ values obtained through simulation follow the same trend as the corresponding theoretical bounds. From Figures 8 and 11 we can observe that the $Max(PoA(SC))$ and $Max(PoA(PX))$ for proportional pricing with polynomial delay are quite small for a wide range of β (from 1 to 10) and n (from 1 to 4). The same can be said about $Max(PoA(SC))$ and $Max(PoA(PX))$ for proportional pricing with queuing delay functions as well, where instead of n, the value of U_f is varied (Figures 9 and 12). Looking closely at $Max(PoA(SC))$ for both type of delay functions we see that if β is chosen to be within a value of 2 to 4, then even the worst case PoA is less than 2. On the other hand, the $Max(PoA(PX))$ for both type of delay functions is quite small (less than 3) for $\beta = 3$ to 5. Since IXPs may want to make good profit while also keeping a low social cost, we observe from our results that $\beta = 3$ or 4 can be a good choice for any IXP to get a good balance between profit and social cost.

5.3.2 *Avg PoA with Polynomial Delay Function*

Simulation results of average PoA (which we denote as $Avg(PoA)$) for proportional and constant pricing with polynomial delay are shown in Figures 10, 13, 14, and 15. To find $Avg(PoA)$, the value of a (recall, $d(y) = ay^n$) that resulted in the worst PoA value for each $\lambda(y)$ curve was considered; the corresponding PoA values were then averaged over all 28 $\lambda(y)$ curves. In other words, $Avg(PoA) = \frac{\sum_{i=1}^{N} \max_a PoA_i}{N}$, where PoA_i is the PoA values found using the i^{th} $\lambda(y)$ curve and N is the total number of $\lambda(y)$ curves, once for each of the $N = 28$ IXPs under consideration.

If we compare Figures 10 and 13 with Figures 8 and 11 respectively, we observe that the $Avg(PoA)$ values are much smaller compared to the their $Max(PoA)$ counterpart for proportional pricing. For $\beta = 3$, the $Avg(PoA(SC))$ has values below 1.1 and $Avg(PoA(PX))$ has values below 1.2, which is quite extraordinary, since it means that with $\beta = 3$, on an average the worst PoA values are well within

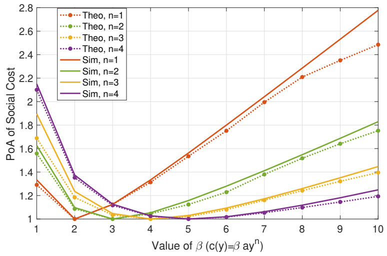

Fig. 8. Max PoA of SC (Sim) with Theoretical bounds (proportional pricing with polynomial delay function).

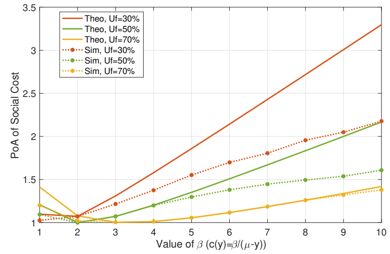

Fig. 9. Max PoA of SC (Sim) with Theoretical bounds (proportional pricing with queuing delay function).

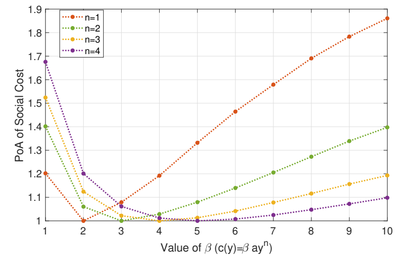

Fig. 10. Avg PoA of SC - Simulated (proportional pricing with polynomial delay).

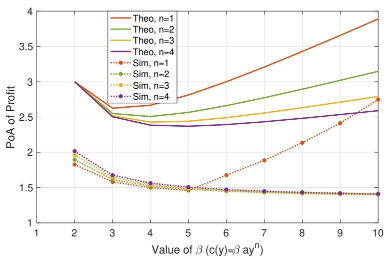

Fig. 11. Max PoA of PX (Sim) with Theoretical bounds (proportional pricing with polynomial delay function).

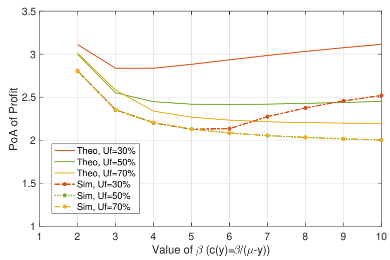

Fig. 12. Max PoA of PX (Sim) with Theoretical bounds (proportional pricing with queuing delay function).

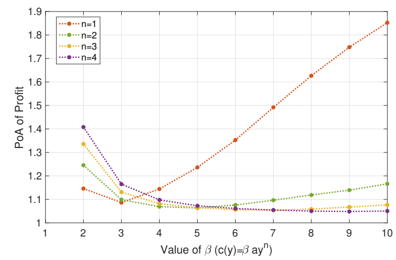

Fig. 13. Avg PoA of PX - Simulated (proportional pricing with polynomial delay).

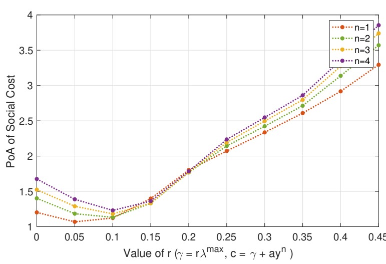

Fig. 14. Avg PoA of SC (Sim) - constant pricing with polynomial delay.

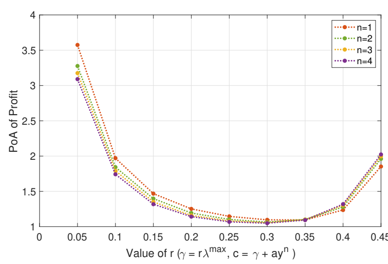

Fig. 15. Avg PoA of PX (Sim) - constant pricing with polynomial delay).

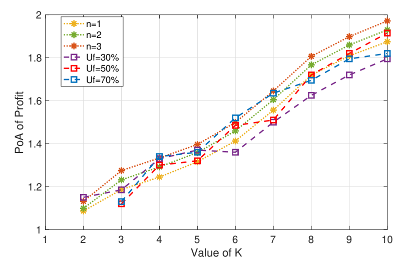

Fig. 16. Avg PoA of PX (Sim) for different values of K (proportional pricing).

20% of the best SC and profit achievable. We see that apart from a linear delay function ($n = 1$), the $Avg(PoA)$ for SC and profit maintain quite a steady small value, and remain small even for higher β values (up to $\beta = 10$). On the other hand, if we look at the results of constant Pricing (Figures 14 and 15), it can be noted that a $\gamma = 0.1\lambda^1$ gives quite a good performance in terms of social cost (having values below 1.25), however the performance degrades heavily with increase in γ. Similar pattern can be noticed for profit as well, where the γ values of $0.15\lambda^1$ to $0.4\lambda^1$ gives pretty reasonable $Avg(PoA(PX))$ values, but it rises on both sides almost exponentially beyond that range. Hence, constant pricing seems less robust and if the constant price γ value is not chosen carefully, then the PoA values for both social cost and profit can be quite bad. Lastly, we see that the value of n does not have much effect on $Avg(PoA)$ values for constant pricing.

The social cost vs profit tradeoff being worse under constant pricing (as compared to proportional pricing) may seem surprising at first thought. This is because given any specific $\lambda(y)$ curve, we can get a constant price that attains the same social cost and IXP profit as proportional pricing. Note however that the results shown here are averaged over different $\lambda(y)$ curves, one for each of the 28 IXPs. This again underlines the fact that when the external routing costs of the participating ISPs are not known to the IXP, proportional pricing are likely to result in better social cost and IXP profit.

5.3.3 Avg PoA with Queuing Delay Function

The $Avg(PoA)$ results with queuing delay functions showed quite similar trend to that of the polynomial delay functions, and due to the space constraint, the results for this case is omitted from plotting. However, we are going to summarize the findings that we observed from simulation. The first thing to notice was that for proportional pricing if we utilize the IXP more (higher U_f), we can gain better PoA results for SC; for profit, the trend seems opposite, at least for lower β values. Also, if we choose $\beta = 3$ or 4, then $Avg(PoA)$ values for both SC and profit turned out to be pretty good, which is consistent to the analysis done in the previous sections.

On the other hand, constant pricing does not seem to have a good trade off between SC and profit with queuing delay function. Increasing the value of γ seemed to make the $Avg(PoA(SC))$ rise sharply, whereas better $Avg(PoA(PX))$ can only be achieved for higher values of γ. Although the effect of U_f is not that prominent on $Avg(PoA(PX))$, but lower U_f results in very bad

$Avg(PoA(SC))$. A reasonable trade-off between SC and profit was obtained when $\gamma = 0.1\lambda^1$, however the results were not very robust to changes in the γ value. Thus, from all the simulation results of constant pricing we can conclude that constant pricing is not very robust, and it is imperative to choose the constant price (γ) precisely to have good performance on both SC and profit. In contrast, proportional pricing yields robust performance over a broad range of β values for both polynomial and queuing delay functions.

5.3.4 Effect of K on $PoA(PX)$

From the theoretical analysis in section 4, we saw that $PoA(PX)$ depends on the value of K, although this dependence is almost linear. We now verify this trend through simulations. Figure 16 show the $Avg(PoA(PX))$ values for proportional pricing (with $\beta = 3$) under polynomial and queuing delay, as a function of K. We observe that the PoA values do increase (almost) linearly with the increment of K, as expected from our theoretical results (Theorems 4.1 and 4.2). However, the PoA values are quite small compared to their corresponding bounds. For both type of delays, even with $K = 10$, the $Avg(PoA(PX))$ have values around 2. While the value of K does depend on the $\lambda(y)$ curve, recalling the definition of K from Section 4.1, K can be expected to be small unless the IXP can carry the traffic demand between all participating IXP's at very low congestion cost. Also, just like our other results, the simulated PoA values are not that sensitive to the value of n and U_f (Fig. 16).

5.3.5 Impact of ISP types and tiers

Note that our analytical and simulation models do not distinguish ISPs based on their tiers or types. ISPs at different tiers or having different types have different geographical footprints, however. For example, an ISP with lower tier number will generally have a larger geographical footprint; the footprint of an access ISP may span a local region or a set of such regions, while a content ISP's footprint may be limited to a few spots where its servers are located. Such footprint and size differences would impact the traffic that they send, and also the potential benefits they may derive from peering at an IXP. In this section, we quantify the traffic sent and/or the costs encountered – both at the IXP and external to the IXP – by ISPs classified according to their tiers (tiers 1, 2, 3) and types (transit, content, and access). The classification of ISP types was directly used from PeeringDB (as discussed in Section 5.1). Since no specific tier-wise classification of ISPs is provided in PeeringDB, we classified the top 7% of the ISPs in the AS rank list as Tier 1, the following 23% as Tier 2 and rest as Tier 3.

Figures 17-22 portrays the amount of traffic, and per unit cost for proportional and zero pricing for different tiers and types of ISPs. Figures 17-19 show results for $\beta = 3$, where the delay function is chosen such that on an average 50% of the traffic is flowing through the IXP and the rest externally. Then for the same delay function, Figures 20-22 show the results when zero pricing is adopted. We observe that costs (per unit traffic) improved across all ISP tiers and types: while zero pricing prompts more ISP pairs to send their traffic through the IXP, it increases the overall cost of traffic due to congestion at the IXP. Proportional pricing motivates some of these ISP pairs to use the external routing option, resulting in lower cost to the ISP - computed as the sum of the external routing cost, the price paid to the IXP, and the congestion (delay) cost at the IXP. These results show that by using proportional pricing, ISPs across different tiers and types all benefit from this cost reduction.

6 Conclusion

We considered the question of pricing of ISP traffic at an IXP, with the goal of attaining an equilibrium solution that is both efficient in terms of social cost and IXP profit. Through both theoretical analysis and simulations, we observed that a pricing policy where the IXP charges the ISPs a per-unit price proportional to the average level of congestion in the public switch, attains a good tradeoff in terms of both of these objectives. A practical benefit of proportional pricing is that a good choice of the price (proportionality constant) does not require the IXP to know the external routing costs of the participating ISPs. This corresponds to the price proportionality factor ($\beta - 1$) being about 2 to 3 times the congestion cost, for which the PoA for both social cost and profit end up being quite small (i.e, less than 2 in general, and often quite close to 1), for both polynomial and queuing congestion cost (delay) functions. Our results also show the performance benefits of proportional pricing (with the proportionality constant chosen appropriately) over a zero pricing scheme where the ISPs experience congestion but no additional price is charged by the IXP. We also analyzed a constant per-unit pricing policy, and argued through theoretical analysis and simulations that the tradeoffs between social cost and IXP profit obtained by constant pricing is not as good as proportional pricing. Furthermore, we observe that choosing a constant price which attains near-optimal social cost and IXP profit requires knowledge of the external routing costs of the IXP. Moreover, the performance under constant pricing is quite sensitive to variations in that constant price, whereas the social cost and IXP profit performance under proportional pricing is more robust to the variations of the price proportionality factor.

In conclusion, we point out a few limitations of our model, addressing which could constitute important directions for future research. Note that our model assumes that the per-unit external routing cost is independent of the traffic amount sent externally. Examining the impact of external route congestion on our proposed policy can be an interesting question for future work. Further, while our analysis includes consideration of asymmetric valuation of traffic between two peering ISPs, it assumes a specific paid peering policy based on Nash bargaining. Exploring the equilibrium efficiency for a broader range of paid peering policies remains open for future investigation.

Appendix A

Towards analyzing the PoA under different pricing schemes, we first state an observation (that will be repeatedly used in our proofs), and prove two useful lemmas that are used in subsequent analysis.

Observation A.1. *For any three numbers, a, b, c, if $a \geq b > c \geq 0$, then $\frac{a-c}{b-c} \geq \frac{a}{b}$.*

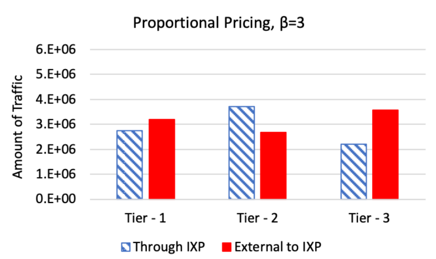

Fig. 17. Traffic flow of ISPs (tier-wise) with proportional pricing.

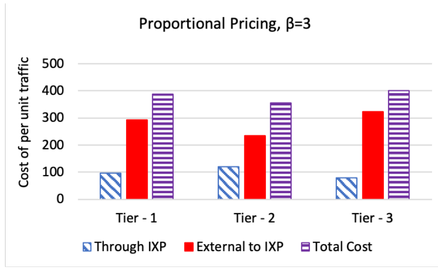

Fig. 18. Per unit cost of traffic (tier-wise) with proportional pricing.

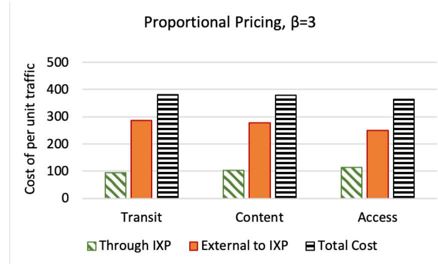

Fig. 19. Per unit cost of traffic (type-wise) with proportional pricing.

Fig. 20. Traffic flow of ISPs (tier-wise) with zero pricing.

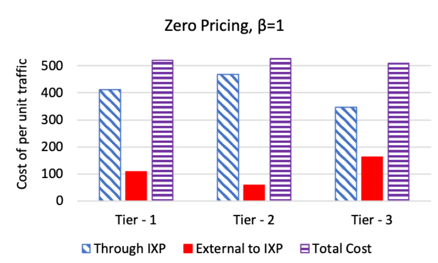

Fig. 21. Per unit cost of traffic (tier-wise) with zero pricing.

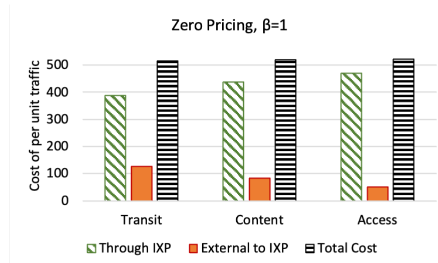

Fig. 22. Per unit cost of traffic (type-wise) with zero pricing.

Lemma A.1. *If $y_e < y_p$, then PoA is bounded by the maximum value of* $\frac{E(y_e)+c(y_e)\cdot(y_p-y_e)}{E(y_p)}$.

Lemma A.2. *If $y_e > y_p$, then PoA is bounded by maximum value of* $\frac{E(y_e)}{E(y_p)+c(y_e)\cdot(y_e-y_p)}$.

Proof of Lemmas A.1 and A.2: For the case of $y_e < y_p$ the PoA equation is:

$$PoA = \max \frac{2\{\int_0^{y_e} E'(y)dy + \int_{y_e}^{y_p} \lambda(y)dy + \int_{y_p}^{y_T} \lambda(y)dy\}}{2\{\int_0^{y_p} E'(y)dy + \int_{y_p}^{y_T} \lambda(y)dy\}}$$

$$\leq \max \frac{\int_0^{y_e} E'(y)dy + \int_{y_e}^{y_p} \lambda(y)dy}{\int_0^{y_p} E'(y)dy} \tag{7}$$

$$\leq \max \frac{E(y_e) + c(y_e)\cdot(y_p - y_e)}{E(y_p)}. \tag{8}$$

The first inequality comes from Observation A.1 and the fact that social cost is minimized at y_p, and the second one comes from the fact that $\lambda(y)$ is non-increasing and within the range $y = y_e$ to $y = y_p$, $\lambda(y)$ can have a maximum value of $\lambda(y_e^+) \leq c(y_e)$.

On the other hand, for the case $y_e > y_p$, we have:

$$PoA = \max \frac{2\{\int_0^{y_e} E'(y)dy + \int_{y_e}^{y_T} \lambda(y)dy\}}{2\{\int_0^{y_p} E'(y)dy + \int_{y_p}^{y_e} \lambda(y)dy + \int_{y_e}^{y_T} \lambda(y)dy\}}$$

$$\leq \max \frac{\int_0^{y_e} E'(y)dy}{\int_0^{y_p} E'(y)dy + \int_{y_p}^{y_e} \lambda(y)dy} \tag{9}$$

$$\leq \max \frac{E(y_e)}{E(y_p) + c(y_e)\cdot(y_e - y_p)}. \tag{10}$$

Similar to the case of $y_e < y_p$, here the first inequality comes from Observation A.1, and the second from $\lambda(y)$ being non-increasing, and that in the range of $y = y_p$ to $y = y_e$, $\lambda(y)$ can have a minimum value of $c(y_e)$. □

Appendix B

B.1 Proofs of Theorem 3.2, Corollary 3.2.1, and 3.2.2

Proof of Theorem 3.2: Corollary 3.2.1 is a special case of Theorem 3.2 case i with $\beta = 1$. Hence proving Theorem 3.2 case i proves Corollary 3.2.1 as well. If $c(y) = \beta\cdot d(y) = \beta\cdot a\cdot y^n$ where $n \geq 1$, then $E'(y) = d(y) + yd'(y) = (n+1)d(y)$; and depending on the value of β, $E'(y)$ or $c(y)$ either can be greater than the other. Comparing both $E'(y)$ and $c(y)$, it is easy to see that if $\beta > n+1$ then $c(y) > E'(y)$ and if $\beta < n+1$ then $c(y) < E'(y)$ for all y. Thus, when $\beta \leq (n+1)$ we use Equation 10 to get,

$$\begin{aligned} PoA &\leq \max \frac{E(y_e)}{E(y_p) + c(y_e)\cdot(y_e - y_p)} \\ &= \max \frac{ay_e^{n+1}}{ay_p^{n+1} + \beta a y_e^n.(y_e - y_p)} \\ &= \max \frac{1}{\rho^{n+1} + \beta - \beta\rho}, \quad where, \rho = y_p/y_e \\ &= \frac{1}{\beta - n(\frac{\beta}{n+1})^{(n+1)/n}}. \end{aligned} \tag{11}$$

The last line is obtained by maximizing the expression in the line above with respect to ρ. Also, if we put $\beta = 1$ in Equation 11, then we get the PoA bound of Corollary 3.2.1. To prove that $\{\frac{1}{\beta - n(\frac{\beta}{n+1})^{(n+1/n)}}\} < (\frac{n+1}{\beta})$, we just have to prove that $\{\beta - n(\frac{\beta}{n+1})^{(n+1/n)}\} > (\frac{\beta}{n+1})$, which can be done by simple algebraic manipulation.

Now, for the case of $\beta > n+1$, we can use Equation 8 to get,

$$\begin{aligned} PoA &\leq \max \frac{E(y_e) + c(y_e)\cdot(y_p - y_e)}{E(y_p)} \\ &= \max \frac{ay_e^{n+1} + \beta a y_e^n(y_p - y_e)}{ay_p^{n+1}} \\ &= \max \{(1-\beta)\rho^{n+1} + \beta\rho^n\}, \quad where, \rho = y_e/y_p \\ &= \frac{\beta}{n+1}[\frac{\beta n}{(\beta-1)(n+1)}]^n. \end{aligned} \tag{12}$$

Proof of Corollary 3.2.2: The proof of Corollary 3.2.2 can be done similarly to those of Corollary 3.2.1 and Theorem 3.2. The only change will be in the value of $d(y)$ that has been used in the derivation. For both the cases of $\beta > n+1$ and $\beta \leq n+1$, $d(y) \leq ay^n$ will be used in the numerator and $d(y) \geq by^n$ will be used in the denominator. The rest of the proof remains the same. □

Proof of Lemma 3.1: To prove Lemma 3.1, we see that it is for the case when $\beta > n+1$. From the argument made in the proof of Theorem 3.2, we know that when $\beta > n+1$, $y_e < y_p$, hence we use Equation 7 to find PoA bound.

$$\begin{aligned}
PoA &\leq \max \ \frac{\int_0^{y_e} E'(y)dy + \int_{y_e}^{y_p} \lambda(y)dy}{\int_0^{y_p} E'(y)dy} \\
&\leq \max \ \frac{\int_{y_e}^{y_p} \lambda(y)dy}{\int_{y_e}^{y_p} E'(y)dy} \leq \max \ \frac{(y_p - y_e)\lambda(y_e)}{(y_p - y_e)E'(y_e)} \\
&= \max \ \frac{\lambda(y_e)}{E'(y_e)} = \max \ \frac{c(y_e)}{E'(y_e)} \\
&= \max \frac{\beta d(y_e)}{d(y_e) + d'(y_e) \cdot y_e} = \frac{\beta}{2}.
\end{aligned}$$

In the second line above, Observation A.1 is used, then in the third line $\lambda(y_e) = c(y_e)$ comes from the fact that the maximum possible value of $\lambda(y_e)$ is $c(y_e)$. Finally, in the last line, the minimum value of $\frac{d'(y_e) \cdot y_e}{d(y_e)}$ is used, which has a minimum value of 1, due to $d(y)$ being convex. □

B.2 Proofs of Theorem 3.4 and Corollary 3.4.1

Theorem 3.4 will be proved for the two cases, when i) $U_f > (1 - \frac{1}{\beta})$ and ii) $U_f < (1 - \frac{1}{\beta})$. Corollary 3.4.1 is a special case of Theorem 3.4, with $U_f > (1 - \frac{1}{\beta})$ and $\beta = 1$, and hence will be proved along with Theorem 3.4.

For proportional pricing we have $c(y) = \beta d(y)$, and $E'(y) = \frac{d}{dy}(\frac{y}{\mu - y}) = \mu d^2(y)$. Let us check what happens to the value of $c(y)$ and $E'(y)$ with increase in y. We see that when no traffic is flowing ($y = 0$), $c(0) = \frac{\beta}{\mu} \geq \frac{1}{\mu} = E'(0)$. Thus, if $\beta = 1$, $c(0) = E'(0)$ and if $\beta > 1$, $c(0) > E'(0)$. Then, if we solve for y such that $c(y) = E'(y)$ and denote it as y_q, we find $y_q = \mu(1 - \frac{1}{\beta})$. Lastly, If $y = y_q + \epsilon$, where $\epsilon > 0$, we find $\frac{E'(y)}{c(y)} = \frac{E'(y_q+\epsilon)}{c(y_q+\epsilon)} = \frac{\mu}{\mu - \beta\epsilon} > 1$. Thus we see that $E'(y)$ starts at a value less than $c(y)$ (if $\beta > 1$), then with the increase of traffic, $E'(y)$ intersects $c(y)$ at traffic y_q, and then stays greater than $c(y)$.

From Theorem 2.1 we know that there is an equilibrium traffic y_e for which $\lambda(y_e^-) \geq c(y_e) \geq \lambda(y_e^+)$. So under different circumstances there might arise the following three scenarios, I) $y_e < y_q$, II) $y_e = y_q$, and III) $y_e > y_q$. For scenario I, $c(y_e) > E'(y_e)$, so $c(y)$ crosses $\lambda(y)$ earlier than $E'(y)$ and hence $y_e \leq y_p$. Arguing similarly, scenario III will have $y_e \geq y_p$. For scenario II, $E'(y) = c(y) = \lambda(y)$ (or, more explicitly $\lambda(y_e^+) \leq E'(y_e) = c(y_e) \leq \lambda(y_e^-)$) and so $y_e = y_p$ (Theorems 2.1 and 2.2). Now we see that, if $y_q = y_e = y_p$ amount of traffic is flowing, then from $y_q = \mu(1 - \frac{1}{\beta})$ we get $U_f = (1 - \frac{1}{\beta})$. However, for scenario I, when $y_e < y_q$, we have $U_f < (1 - \frac{1}{\beta})$, and for scenario III of $y_q < y_e$, we have $U_f > (1 - \frac{1}{\beta})$.

For the case $U_f > 1 - 1/\beta$, we see that the condition of Lemma A.2 is satisfied, and hence from Equation 10,

$$\begin{aligned}
PoA &\leq \max \ \frac{E(y_e)}{E(y_p) + c(y_e) \cdot (y_e - y_p)} \\
&= \max \ \frac{\frac{y_e}{\mu - y_e}}{\frac{y_p}{\mu - y_p} + \frac{\beta}{\mu - y_e} \cdot (y_e - y_p)} \\
&= \max \frac{U_f(1-\rho)}{\rho(1-U_f) + \beta(1-\rho)(U_f - \rho)}. \qquad (13)
\end{aligned}$$

In the above calculation, the third line comes from multiplying both numerator and denominator with $\frac{(\mu - y_e)(\mu - y_p)}{\mu^2}$, and denoting $y_p/\mu = \rho$. As we can see from Equation 13, the maximum value of PoA depends on the value of $\rho = y_p/\mu$ and $U_f = y_e/\mu$. Thus, for a given U_f if we can find the value of ρ which gives the maximum PoA, and use that to get the PoA upper bound from Equation 13. This yields, $\rho = 1 - \sqrt{\frac{1-U_f}{\beta}}$. Now putting this value of ρ in Equation 13 we get the PoA bound of Theorem 3.4 case i, which is:

$$PoA \leq \frac{U_f \sqrt{\frac{1-U_f}{\beta}}}{(1-U_f)[2 - \sqrt{1-U_f}(\frac{1+\beta}{\sqrt{\beta}})]}. \qquad (14)$$

Now, if we put $\beta = 1$ in the above equation, we get the PoA bound of Corollary 3.4.1 as well. On the other hand, for the case of $U_f < 1 - 1/\beta$, since the condition of Lemma A.1 is satisfied, Equation 8 can be used to find the maximum value of PoA.

$$\begin{aligned}
PoA &\leq \max \ \frac{E(y_e) + c(y_e) \cdot (y_p - y_e)}{E(y_p)} \\
&= \max \ \frac{\frac{y_e}{\mu - y_e} + \frac{\beta}{\mu - y_e}(y_p - y_e)}{\frac{y_p}{\mu - y_p}} \\
&= \max \ \frac{1-\rho}{\rho} \times \frac{U_f(1-\beta) + \beta\rho}{1 - U_f} \\
&= \frac{\{\sqrt{\beta} - \sqrt{U_f(\beta - 1)}\}^2}{1 - U_f}. \qquad (15)
\end{aligned}$$

In the above calculation, the third line comes from multiplying both numerator and denominator with $\frac{(\mu - y_e)(\mu - y_p)}{\mu^2}$ and denoting $y_p/\mu = \rho$. The expression in the last line is obtained by plugging in the value of ρ, which maximizes the expression before it, for a given U_f. □

Appendix C

C.1 Proof of Theorem 4.1

To get the PoA of profit we need to use the fact that maximum achievable profit ($max(PX)$) is at most equal to the maximum social welfare ($SW(opt)$). To prove that, let us define social welfare (SW) in terms of social cost.

Definition C.1. *Social welfare is equal to total utility achievable minus the social cost.*

Thus, from the above definition, for our case $SW = \int_0^{y_T} \lambda(y)dy - SC$, and the lower the value of SC the higher will be the value of SW. Also, we know SC is minimum at a traffic y_p and hence $SW(opt)$ is achieved at traffic y_p as well. Now, let us say y_r be the traffic which gives the maximum profit. Then profit will be $(\beta - 1)d(y_r) \cdot y_r$ or $\gamma \cdot y_r$, under proportional and constant pricing respectively. So, for proportional pricing with polynomial delay (with $y_r = y_e \leq y_p$) we have: $\frac{max(PX)}{SW(opt)}$

$$= \frac{(\beta-1)d(y_r)\cdot y_r}{\int_0^{y_p}\lambda(y)-\int_0^{y_p}E'(y)dy} \le \frac{(\beta-1)d(y_r)\cdot y_r}{\int_0^{y_r}\lambda(y)-\int_0^{y_r}E'(y)dy}$$
$$\le \frac{(\beta-1)E(y_r)}{y_r\cdot c(y_r)-E(y_r)} \le \frac{(\beta-1)E(y_r)}{\beta E(y_r)-E(y_r)} \le 1.$$

In a similar way, we can prove that for $y_r = y_e \ge y_p$ the same inequality holds. Moreover, the proofs can be extended to constant pricing as well. In all the following bound calculations for $PoA(PX)$, the fact that the maximum achievable profit ($\max(PX)$) is at most equal to maximum social welfare ($SW(opt)$) will be used.

Similar to the PoA calculation of social cost as in previous sections, PoA of profit need to be calculated for two cases, i) $y_e \ge y_p$ or equivalently, $\beta \ge n+1$ and ii) $y_e < y_p$ or equivalently, $\beta < n+1$. Hence, from the definition of price of anarchy we get PoA for $y_e \ge y_p$ as,

$$\begin{aligned}
PoA(PX) &= \max \frac{Maximum\ Achievable\ Profit}{Profit\ at\ Equilibrium} \\
&\le \max \frac{Optimum\ Social\ Welfare}{Profit\ at\ Equilibrium} \\
&\le \max \frac{\int_0^{y_p}\lambda(y)-\int_0^{y_p}E'(y)}{c(y_e)\cdot y_e-d(y_e)\cdot y_e} \\
&\le \max \frac{\lambda^1 y_e+(y_p-y_e)c(y_e)-E(y_p)}{(\beta-1)ay_e^{n+1}} \\
&\le \max \frac{Kc(y_e)y_e+(y_p-y_e)c(y_e)-E(y_p)}{(\beta-1)ay_e^{n+1}}, [\lambda^1 \le K\lambda_e] \\
&\le \max \frac{K\beta ay_e^{n+1}+(y_p-y_e)\beta ay_e^n-ay_p^{n+1}}{(\beta-1)ay_e^{n+1}} \\
&\le \max \frac{K\beta+(\rho-1)\beta-\rho^{n+1}}{\beta-1} \text{ [asm. } \rho = y_p/y_e] \\
&= \frac{K\beta+((\frac{\beta}{n+1})^{1/n}-1)\beta-(\frac{\beta}{n+1})^{(n+1)/n}}{\beta-1},
\end{aligned} \tag{16}$$

where the expression in the last line is obtained by maximizing the expression before it with respect to ρ. Now, (in a similar way) for the case when $y_e < y_p$ we have,

$$\begin{aligned}
PoA(PX) &\le \max \frac{\int_0^{y_p}\lambda(y)-\int_0^{y_p}E'(y)}{c(y_e)\cdot y_e-d(y_e)\cdot y_e} \\
&\le \max \frac{\lambda^1 y_p-E(y_p)}{(\beta-1)ay_e^{n+1}} \\
&\le \max \frac{Kc(y_e)y_p-E(y_p)}{(\beta-1)ay_e^{n+1}} \text{ [asm. } \lambda^1 = K\lambda_e] \\
&\le \max \frac{K\beta ay_e^n y_p-ay_p^{n+1}}{(\beta-1)ay_e^{n+1}} \\
&\le \max \frac{K\beta\rho-\rho^{n+1}}{\beta-1} \quad \text{[assuming } \rho = y_p/y_e].
\end{aligned}$$

Note that $\rho = (\frac{K\beta}{n+1})^{1/n}$ maximizes the term in the right hand of the final inequality above. Since $\max(\rho) = 1$, the value of ρ for maximum PoA is $\min((\frac{K\beta}{n+1})^{1/n}, 1)$. □

C.2 Proof of Theorem 4.2

Similar to the PoA analysis for SC, we determine PoA bounds of profit for the case of, i) $U_f \le 1-1/\beta$ and ii) $U_f \ge 1-1/\beta$. When $U_f \le 1-1/\beta$, we have:

$$\begin{aligned}
PoA(PX) &\le \max \frac{Optimum\ Social\ Welfare}{Profit\ at\ Equilibrium} \\
&\le \max \frac{\int_0^{y_p}\lambda(y)-\int_0^{y_p}E'(y)}{c(y_e)\cdot y_e-d(y_e)\cdot y_e} \\
&\le \max \frac{\lambda^1 y_e+(y_p-y_e)c(y_e)-E(y_p)}{\frac{\beta-1}{\mu-y_e}y_e} \\
&\le \max \frac{Kc(y_e)y_e+(y_p-y_e)c(y_e)-E(y_p)}{\frac{\beta-1}{\mu-y_e}y_e}, [\lambda^1 = K\lambda_e] \\
&\le \max \frac{K\frac{\beta y_e}{\mu-y_e}+(y_p-y_e)\frac{\beta}{\mu-y_e}-\frac{y_p}{\mu-y_p}}{\frac{\beta-1}{\mu-y_e}\cdot y_e} \\
&\le \max \frac{K\beta U_f(1-\rho)+\beta(\rho-U_f)(1-\rho)-\rho(1-U_f)}{(\beta-1)(1-\rho)U_f} \\
&= \frac{\beta U_f(K-1)+(\sqrt{1-U_f}-\sqrt{\beta})^2}{U_f(\beta-1)},
\end{aligned} \tag{17}$$

where in the second line from the last, both numerator and denominator were multiplied with $\frac{(\mu-y_e)(\mu-y_p)}{\mu^2}$, and the last line is obtained by maximizing the expression above it with respect to ρ. On the other hand, for the case of $U_f \ge 1-1/\beta$, we have: $PoA(PX)$

$$\begin{aligned}
&\le \max \frac{\int_0^{y_p}\lambda(y)-\int_0^{y_p}E'(y)}{c(y_e)\cdot y_e-d(y_e)\cdot y_e} \le \max \frac{\lambda^1 y_p-E(y_p)}{\frac{\beta-1}{\mu-y_e}y_e} \\
&\le \max \frac{Kc(y_e)y_p-E(y_p)}{\frac{\beta-1}{\mu-y_e}y_e} \quad \text{[assuming } \lambda^1 = K\lambda_e] \\
&\le \max \frac{K\frac{\beta y_p}{\mu-y_e}-\frac{y_p}{\mu-y_p}}{\frac{\beta-1}{\mu-y_e}\cdot y_e} \le \max \frac{K\beta y_p(\mu-y_p)-y_p(\mu-y_e)}{(\beta-1)(\mu-y_p)y_e} \\
&\le \max \frac{K\beta\rho(1-\rho)-\rho(1-U_f)}{(\beta-1)(1-\rho)U_f} = \frac{(\sqrt{K\beta}-\sqrt{1-U_f})^2}{U_f(\beta-1)},
\end{aligned}$$

where in the second line from the last, we multiplied both numerator and denominator with $(\mu-y_e)(\mu-y_p)$. In the last line both numerator and denominator were divided by μ^2, and the final expression is obtained by plugging in the value of ρ that maximizes the expression just before it. □

Acknowledgment

This work was supported in part by the National Science Foundation through awards CNS-1816396 and CNS-1814086.

References

[1] M. I. Alam, E. Anshelevich, K. Kar, and M. Yuksel, "Proportional pricing for efficient traffic equilibrium at internet exchange points," in *2021 33th International Teletraffic Congress (ITC-33)*. IEEE, 2021, pp. 1–9.

[2] M. Chiesa, D. Demmler, M. Canini, M. Schapira, and T. Schneider, "SIXPACK: Securing Internet EXchange Points Against Curious OnlooKers," in *Proceedings of the 13th International Conference on Emerging Networking EXperiments and Technologies (CoNEXT '17)*. New York, NY, USA: ACM, 2017, p. 120–133.

[3] B. Ager, N. Chatzis, A. Feldmann, N. Sarrar, S. Uhlig, and W. Willinger, "Anatomy of a Large European IXP," in *Proceedings of the ACM SIGCOMM 2012 conference on Applications, technologies, architectures, and protocols for computer communication*, 2012, pp. 163–174.

[4] EuroIX, "European Internet Exchange Association (EuroIX) - Reports," 2021, retrieved on April 03, 2023. [Online]. Available: https://www.euro-ix.net/en/services/euro-ix-reports/

[5] P. S. Ryan and J. Gerson, "A Primer on Internet Exchange Points for Policymakers and Non-Engineers," http://ssrn.com/abstract=2128103, August 2012.

[6] Equinix, "Equinix Internet Exchange," January 2021, retrieved on January 18, 2023. [Online]. Available: https://www.equinix.com/interconnection-services/ internet-exchange/

[7] W. Norton, *The Internet Peering Playbook: Connecting to the Core of the Internet*. DrPeering Press, 2012.

[8] K. Russell, "Here's Why You Should Be Thrilled Netflix Is Paying Comcast For Content Delivery," Business Insider, Feb. 25, 2014.

[9] P. K. Dey and M. Yuksel, "Peering Among Content-Dominated Vertical ISPs," *IEEE Networking Letters*, vol. 1, no. 3, pp. 132–135, 2019.

[10] J. C. Cardona Restrepo and R. Stanojevic, "IXP Traffic: a Macroscopic View," in *Proceedings of the 7th Latin American Networking Conference*, 2012, pp. 1–8.

[11] A. Dhamdhere and C. Dovrolis, "The Internet is flat: modeling the transition from a transit hierarchy to a peering mesh," in *Proceedings of ACM CoNEXT*, 2010, pp. 1–12.

[12] C. Labovitz, S. Iekel-Johnson, D. McPherson, J. Oberheide, and F. Jahanian, "Internet Inter-Domain Traffic," *ACM SIGCOMM Computer Communication Review*, vol. 40, no. 4, pp. 75–86, 2010.

[13] T. Böttger, G. Antichi, E. L. Fernandes, R. di Lallo, M. Bruyere, S. Uhlig, and I. Castro, "The Elusive Internet Flattening: 10 Years of IXP Growth," *arXiv e-prints*, 2018.

[14] M. S. Schwartz and E. Mershon, "Paid Internet Peering on the Rise, Disputes Possible," *Communications Daily*, vol. 33, no. 126, 2013.

[15] X. Wang, Y. Xu, and R. T. B. Ma, "Paid Peering, Settlement-Free Peering, or Both?" in *IEEE INFOCOM 2018 - IEEE Conference on Computer Communications*, 2018, pp. 2564–2572.

[16] N. Chatzis, G. Smaragdakis, A. Feldmann, and W. Willinger, "There is More to IXPs than Meets the Eye," *SIGCOMM Comput. Commun. Rev.*, vol. 43, no. 5, p. 19–28, Nov. 2013.

[17] N. Chatzis, G. Smaragdakis, J. Böttger, T. Krenc, and A. Feldmann, "On the Benefits of Using a Large IXP As an Internet Vantage Point," in *Proceedings of ACM Internet Measurement Conference (IMC)*. New York, NY, USA: ACM, 2013, pp. 333–346.

[18] M. V. B. da Silva, A. S. Jacobs, R. J. Pfitscher, and L. Z. Granville, "IDEAFIX: identifying elephant flows in P4-based IXP networks," in *Proceedings of IEEE Global Communications Conference (GLOBECOM)*, 2018, pp. 1–6.

[19] K.-K. Yap, M. Motiwala, J. Rahe, S. Padgett, M. Holliman, G. Baldus, M. Hines, T. Kim, A. Narayanan, A. Jain, V. Lin, C. Rice, B. Rogan, A. Singh, B. Tanaka, M. Verma, P. Sood, M. Tariq, M. Tierney, D. Trumic, V. Valancius, C. Ying, M. Kallahalla, B. Koley, and A. Vahdat, "Taking the edge off with espresso: Scale, reliability and programmability for global internet peering," in *Proceedings of the Conference of the ACM Special Interest Group on Data Communication*, 2017, p. 432–445.

[20] M. Bruyere, R. Lapeyrade, E. L. Fernandes, I. Castro, S. Uhlig, A. W. Moore, and G. Antichi, "Umbrella: A deployable SDN-enabled IXP switching fabric," in *Proceedings of the Symposium on SDN Research*, 2018.

[21] A. Gupta, L. Vanbever, M. Shahbaz, S. P. Donovan, B. Schlinker, N. Feamster, J. Rexford, S. Shenker, R. Clark, and E. Katz-Bassett, "SDX: A software defined internet exchange," in *Proceedings of ACM SIGCOMM*, 2014, pp. 551–562.

[22] M. Bruyere, G. Antichi, E. L. Fernandes, R. Lapeyrade, S. Uhlig, P. Owezarski, A. W. Moore, and I. Castro, "Rethinking IXPs' architecture in the age of SDN," *IEEE Journal on Selected Areas in Communications*, vol. 36, no. 12, pp. 2667–2674, 2018.

[23] G. Antichi, I. Castro, M. Chiesa, E. L. Fernandes, R. Lapeyrade, D. Kopp, J. H. Han, M. Bruyere, C. Dietzel, M. Gusat, A. W. Moore, P. Owezarski, S. Uhlig, and M. Canini, "ENDEAVOUR: a scalable SDN architecture for real-world IXPs," *IEEE Journal on Selected Areas in Communications*, vol. 35, no. 11, pp. 2553–2562, 2017.

[24] LINX, "LINX: London Internet Exchange Service Fees," retrieved on November 6, 2023. [Online]. Available: https://www.linx.net/services/service-fees/

[25] Equinix, "Equinix: What is colocation?" retrieved on November 6, 2023. [Online]. Available: https://www.equinix.com/what-is-colocation

[26] AMS-IX, "Amsterdam Internet Exchange: Pricing," retrieved on November 6, 2023. [Online]. Available: https://www.ams-ix.net/ams/pricing

[27] Nisan, Noam and Roughgarden, Tim and Tardos, Eva and Vazirani, Vijay V., *Algorithmic Game Theory*. USA: Cambridge University Press, 2007.

[28] T. Roughgarden, "Intrinsic robustness of the price of anarchy," *Journal of the ACM (JACM)*, vol. 62, no. 5, pp. 1–42, 2015.

[29] E. Tardos and T. Wexler, "Network formation games and the potential function method," *Algorithmic Game Theory*, pp. 487–516, 2007.

[30] M. O. Jackson and A. Wolinsky, "A strategic model of social and economic networks," *Journal of economic theory*, vol. 71, no. 1, pp. 44–74, 1996.

[31] V. Bala and S. Goyal, "A noncooperative model of network formation," *Econometrica*, vol. 68, no. 5, pp. 1181–1229, 2000.

[32] J. Derks, J. Kuipers, M. Tennekes, and F. Thuijsman, "Local dynamics in network formation," in *Proc. Third World Congress of The Game Theory Society*. Citeseer, 2008.

[33] B. Dutta, S. Ghosal, and D. Ray, "Farsighted network formation," *Journal of Economic Theory*, vol. 122, no. 2, pp. 143–164, 2005.

[34] M. O. Jackson, "A Survey of Network Formation Models: Stability and Efficiency," *Group formation in economics: Networks, clubs, and coalitions*, vol. 664, pp. 11–49, 2005.

[35] A. Fabrikant, A. Luthra, E. Maneva, C. H. Papadimitriou, and S. Shenker, "On a Network Creation Game," in *Proceedings of the twenty-second annual symposium on Principles of distributed computing*, 2003, pp. 347–351.

[36] J. Corbo and D. Parkes, "The Price of Selfish Behavior in Bilateral Network Formation," in *Proceedings of the twenty-fourth annual ACM symposium on Principles of distributed computing*, 2005, pp. 99–107.

[37] F. Page and J. Resende, "Network formation games," http://kenticoqa.uwb.edu/getattachment/business/about/research-series/network-formation-games-uwb.pdf, 2013.

[38] A. Calvó-Armengol and R. İlkılıç, "Pairwise-stability and nash equilibria in network formation," *International Journal of Game Theory*, vol. 38, no. 1, pp. 51–79, 2009.

[39] E. Anshelevich and M. Hoefer, "Contribution games in networks," *Algorithmica*, vol. 63, no. 1-2, pp. 51–90, 2012.

[40] S. Shakkottai and R. Srikant, "Economics of network pricing with multiple isps," *IEEE/ACM Transactions on Networking*, vol. 14, no. 6, pp. 1233–1245, 2006.

[41] C.-K. Chau, Q. Wang, and D.-M. Chiu, "On the viability of paris metro pricing for communication and service networks," in *2010 Proceedings IEEE INFOCOM*. IEEE, 2010, pp. 1–9.

[42] P. Hande, M. Chiang, R. Calderbank, and J. Zhang, "Pricing under constraints in access networks: Revenue maximization and congestion management," in *2010 Proceedings IEEE INFOCOM*. IEEE, 2010, pp. 1–9.

[43] R. T. Ma, J. C. Lui, and V. Misra, "Evolution of the internet economic ecosystem," *IEEE/ACM Transactions on Networking*, vol. 23, no. 1, pp. 85–98, 2013.

[44] V. Valancius, C. Lumezanu, N. Feamster, R. Johari, and V. V. Vazirani, "How Many Tiers? Pricing in the Internet Transit Market," *ACM SIGCOMM Computer communication review*, vol. 41, no. 4, pp. 194–205, 2011.

[45] R. T. B. Ma, D. M. Chiu, J. C. S. Lui, V. Misra, and D. Rubenstein, "On cooperative settlement between content, transit, and eyeball internet service providers," *IEEE/ACM Transactions on Networking*, vol. 19, no. 3, pp. 802–815, 2011.

[46] E. Anshelevich, O. Bhardwaj, and K. Kar, "Strategic Network Formation Through an Intermediary," *Theory of Computing Systems*, vol. 63, no. 6, pp. 1314–1335, 2019.

[47] D. Peering, "A Brief Study of 28 Peering Policies," retrieved on April 03, 2023. [Online]. Available: http://drpeering.net/AskDrPeering/blog/articles/

[48] PeeringDB, "PeeringDB Website," retrieved on April 03, 2023. [Online]. Available: https://www.peeringdb.com

[49] D. Zarchy, A. Dhamdhere, C. Dovrolis, and M. Schapira, "Nash-Peering: A New Techno-Economic Framework for Internet Interconnections," in *IEEE INFOCOM 2018-IEEE Conference on Computer Communications Workshops (INFOCOM WKSHPS)*. IEEE, 2018, pp. 403–408.

[50] Y. Zhang, M. Roughan, C. Lund, and D. L. Donoho, "Estimating Point-to-Point and Point-to-Multipoint Traffic Matrices: an Information-Theoretic Approach," *IEEE/ACM Transactions on Networking*, vol. 13, no. 5, pp. 947–960, October 2005.